Huygens Essential - Visualization and Analysis

Visualization and Analysis User Guide for version 3.7

Scientific Volume Imaging B.V.

Huygens Essential - Visualization and Analysis

Visualization and Analysis User Guide for version 3.7

Scientific Volume Imaging B.V.

Copyright © 1995-2011 by Scientific Volume Imaging B.V.
All rights reserved

Mailing Address	Scientific Volume Imaging B.V. Laapersveld 63 1213 VB Hilversum The Netherlands
Phone	+31 35 6421626
Fax	+31 35 6837971
E-mail	info@svi.nl
URL	http://www.svi.nl/

Cover illustration: Macrophage recorded by Dr. James Evans (White-head Institute, MIT, Boston MA, USA) using widefield microscopy, as deconvolved with Huygens®. Stained for tubulin (yellow/green), actin (red) and the nucleus (DAPI, blue).

Contents

CHAPTER 1 *Introduction* *1*

CHAPTER 2 *Installation* *3*

Microsoft Windows 3
Microsoft Windows 64 bit Edition 3
Mac OS X 3
Linux (Debian) 3
Linux (RPM) 4
After the Installation 4
The License String 4
Updating the Software 6
Removing the Software 7
System Requirements for Huygens Essential 7
Support on Installation 8

CHAPTER 3 *The Image Restoration Process* *9*

The Processing Stages in the Wizard 9
Loading an Image 9
Verifying Microscopic Parameters 10
The Intelligent Cropper 12
The Image Histogram 13
Estimating the Average Background 14
The Deconvolution Stage 15
Finishing or Restarting a Deconvolution Run 16
Multi-channel Images 16
Z-drift Correcting for Time Series 17
Saving the Result 17
Using a Measured PSF 17

CHAPTER 4 *The Twin Slicer* ... 19
Using the Slicer in Basic Mode ... 20
Using the Slicer in Advanced Mode ... 22
Measurement ... 24

CHAPTER 5 *The Orthogonal Slicer* ... 27
The Crosshair Cursor ... 28
Visualization parameters ... 28
Measurements ... 28
Auto-Zoom ... 28
Display Options ... 28

CHAPTER 6 *The MIP Renderer* ... 29
Basic Usage ... 29
Advanced Usage ... 31
Simple Animations ... 31

CHAPTER 7 *The SFP Renderer* ... 33
Basic Usage ... 33
Advanced Usage ... 35
Simple Animations ... 36

CHAPTER 8 *The Surface Renderer* ... 37
Basic Usage ... 38
Advanced Usage ... 39
Simple Animations ... 40

CHAPTER 9 *The Movie Maker* ... 43
An Overview ... 43
Creating and Adjusting Keyframes ... 44
Using the Storyboard ... 45
Working with Movie Projects ... 46
Using the Timeline ... 46
Advanced Topics ... 47

CHAPTER 10 *Introduction to the Object Analyzer* ... 49
Starting the Object Analyzer ... 49
Segmenting the Objects: Setting the Threshold ... 51
Interaction with the Objects ... 53

Render Pipes 54
Object Statistics 55
Storing your Results 59
Further Reading 59

CHAPTER 11 *Object Analyzer Geometry Measurements* *61*

Iso-surface 61
Principal Axis 61
Length and width 61
Sphericity 62
Aspect Ratio 63
More Parameters and Filtering 63

CHAPTER 12 *Object Analyzer Component Reference* *65*

Main window components 66

CHAPTER 13 *The Colocalization Analyzer* *81*

How to use the Colocalization Analyzer 81
Iso-colocalization object analysis 84
Backgrounds vs. thresholds in colocalization 84

CHAPTER 14 *The Chromatic Shift Corrector* *87*

Starting the Chromatic Shift Corrector 87
Estimation of the chromatic shifts 88
Visualization of the chromatic shifts 88
Editing the chromatic shift vectors 89
Working with templates 90

CHAPTER 15 *Huygens Remote Manager* *91*

What is HRM 91
Where to find HRM 92
HRM requirements and technical features 92
How HRM communicates with Huygens Core 93
Basic HRM usage 94

CHAPTER 16 *Appendix* *101*

License String Details 101
The Point Spread Function 102
Quality Factor 103

File Series 103
Adjusting the Global Color Scheme 104
Hue Selector 105
Image Statistics 105
Setting the Coverslip Position 106
Excitation Beam Overfill Factor 107
Brightfield Images 108
Support and Contact Information 108

CHAPTER 1

Introduction

Huygens Essential is an image processing software package tailored for restoration, visualization and analysis of microscopic images. Its wizard driven user interface guides through the process of deconvolving images from light microscopes. Huygens Essential is able to deconvolve a wide variety of images ranging from 2D widefield images to 4D multi-channel multi-photon confocal time series. To facilitate comparison of raw and deconvolved data or results from different deconvolution runs Huygens Essential is equipped with a dual 4D slicer tool. Also 3D images and animations can be rendered with its powerful visualization tools. Post-restoration analysis is possible using the interactive analysis tools.

Based on the same image processing engine (the compute engine) as Huygens Professional, Huygens Essential combines the quality and speed of the algorithms available in Huygens Professional with the ease of use of a wizard driven intelligent user interface fortified with a versatile and intuitive batch processor.

Huygens Essential uses cross-platform technology. It is available on Microsoft Windows 2003 Server, XP (32 and 64 bit), Vista (32 and 64 bit), and Windows 7 (32 and 64 bit), Linux (32 and 64 bit), and Mac OS X Tiger (32 bit only) and (Snow)Leopard (32 and 64 bit). IRIX and Itanium distributions are available on demand.

CHAPTER 2

Installation

Huygens Essential can be downloaded from the SVI website[1].

Microsoft Windows

Double click on the Huygens installer executable, e.g. `huygens-370p0.exe`. Double click its icon to start the installation. During installation the directory `C:\Program files\SVI\` will be created by default. After completion the four Huygens icons appear on the desktop. Double clicking on the Huygens Essential icon starts the program.

Microsoft Windows 64 bit Edition

Double click the Huygens installer executable, e.g. `huygens-370p0_x86_64.exe`. Note that the 64 bit Windows version will only run on 64 bit editions of Microsoft Windows 7, Vista and XP. During installation the directory `C:\Program files (x86)\SVI\` will be created by default. Both the 32 and 64 bit Huygens versions will be installed in this directory. After completion the four Huygens icons appear on the desktop.

Mac OS X

Double click the package file, for instance `huygens-3.7.0-p0-Leopard-i386.pkg.tar.gz`. The archive manager expands it to a `.pkg` file, which will be placed in the same directory. Double click this file, and follow the installation wizard.

Linux (Debian)

Debian packages are natively used by Ubuntu and other Debian-based Linux distributions. Double click the package file, .e.g. `huygens-3.7.0-p0_i386.deb`, and follow the steps in the package manager. To install the package through the command line:

```
dpkg -i huygens-3.7.0-p0_i386.deb
```

1. `http://www.svi.nl/`

Linux (RPM)

RPM (RedHat Package Manager) packages are natively used by RedHat, Fedora, SUSE, and other RPM-based Linux distributions. Double click the package file, e.g. `huygens-3.7.0-p0.i386.rpm`, and follow the steps in the package manager, or install the package through the command line:

```
rpm -ivh --force huygens-3.7.0-p0.i386.rpm
```

After the Installation

After a first-time installation there is not yet a license available. However, still the software can be start ed. Without a license it will run in *Freeware mode*. The *System ID*, necessary for generating a license, is obtained by pressing the GET A LICENSE button (See Figure 2.1) when opening Huygens Essential and it can be found in the HELP→LICENSE menu. The next section explains how to obtain and install a license string.

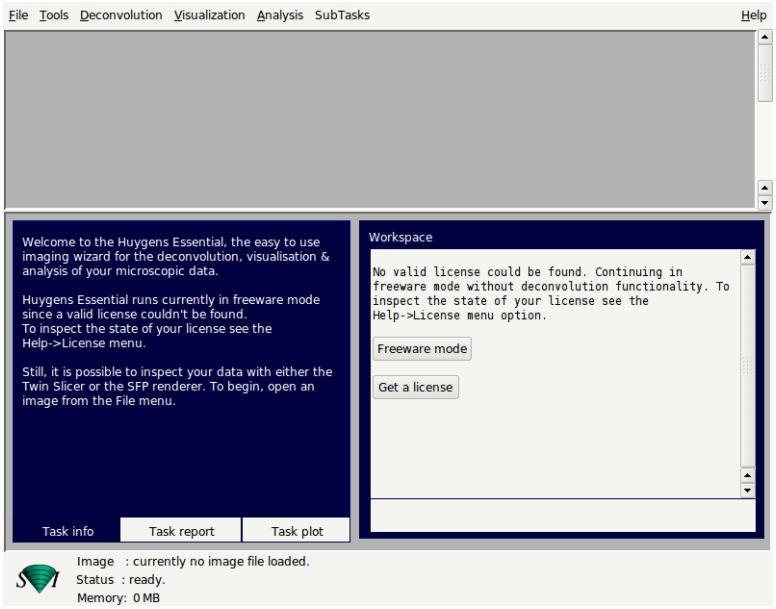
File Tools Deconvolution Visualization Analysis SubTasks Help
Welcome to the Huygens Essential, the easy to use imaging wizard for the deconvolution, visualisation & analysis of your microscopic data.
Huygens Essential runs currently in freeware mode since a valid license couldn't be found.
To inspect the state of your license see the Help->License menu.
Still, it is possible to inspect your data with either the Twin Slicer or the SFP renderer. To begin, open an image from the File menu.
Task info
Task report
Task plot
Workspace
No valid license could be found. Continuing in freeware mode without deconvolution functionality. To inspect the state of your license see the Help->License menu option.
Freeware mode
Get a license
Image : currently no image file loaded.
Status : ready.
Memory: 0 MB

FIGURE 2.1. The startup window of Huygens Essential. If no license string is installed the software runs in *freeware mode*.

The License String

The license key used by all SVI software is a single string per licensed package. It may look as follows:

```
HuEss-3.7-wcnp-d-tvAC-emnps-eom2012Dec31-
e7b7c623393d708e-{user@domain.com}-4fce0dbe86e8ca4344dd
```

At startup Huygens Essential searches for a license file `huygensLicense` which contains a license string. This license string is provided by SVI via e-mail. Installing the license string is the same for all platforms.

Obtaining a License String

If upgrading is not handled from a previous installation it is likely that a license is not yet available. To enable us to generate a license string, we need the *fingerprint* of the computer used, the so called *system ID* number. If Huygens Essential is not already running, please start it. The system ID pops up as long as no valid license is available and is displayed in the HELP→ABOUT dialog (Figure 2.2). Send it to `sales@svi.nl`, and a license string will be provided. To prevent any typing error use the COPY button to save the ID to the clipboard. It can be printed into the license mail message with the EDIT→PASTE menu item of the mail program.

This dialog box also contains a button to *Check for updates* on the SVI company server.

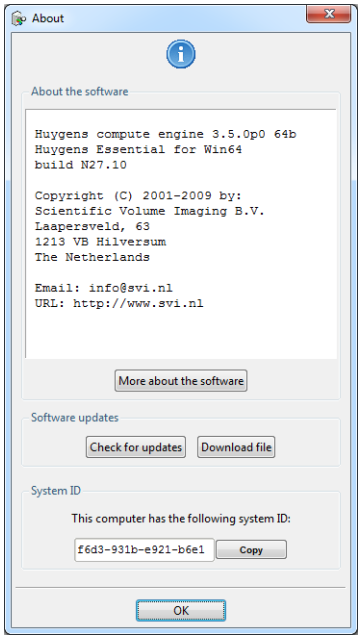

FIGURE 2.2. The HELP→ABOUT window. The system ID is shown at the bottom.

Installing the License String

Select the license string in the e-mail message and copy it to the clipboard using EDIT→COPY in the mailing program. Start Huygens Essential and go to HELP→LICENSE: a dialog box pops up. Then press the ADD NEW LICENSE button and paste the string into the text field (Figure 2.3). Complete the procedure by pressing ADD LICENSE; this will

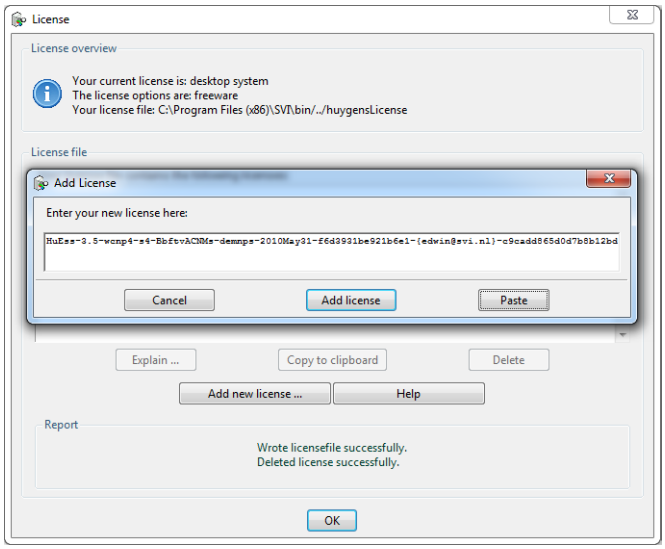

FIGURE 2.3. The license window allows to add, delete and troubleshoot licenses.

add the string to the `huygensLicense` file. Please try to avoid typing the license string by hand: any typing error will invalidate the license. With an invalid license, the software will remain in Freeware mode. When the license is correct the message "*Added license successfully*" will appear.
Restart Huygens Essential to activate the new license.

Location of the License File

The license string is added to the file huygensLicense in the SVI directory (Table 2.1 on page 6).

TABLE 2.1. The default installation paths per platform.

Platform	Installation path
Windows	`C:\Program files\SVI\`
Windows 64 bit Edition	`C:\Program files (x86)\SVI\`
Mac OS X	`/Applications/SVI/`[a]
Linux	`/usr/local/svi/`

a. The path name on Mac OS X depends on where the software is installed. This is a typical example.

On Irix and Linux and Mac OS X an alternative location is the user's home directory. On OS X this is especially convenient when updating frequently.

Troubleshooting License Strings

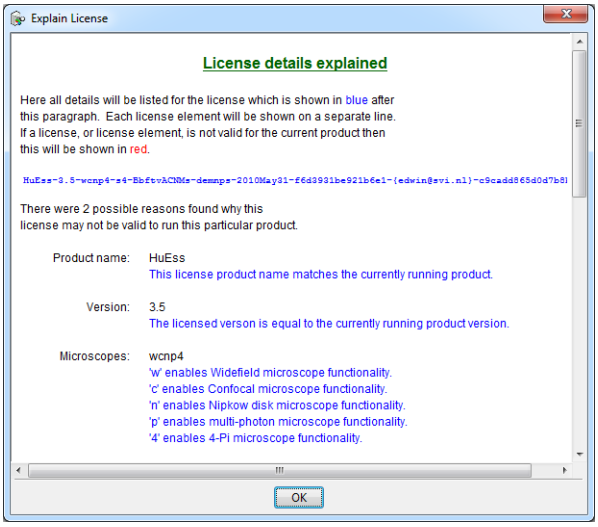
Explain License
License details explained
Here all details will be listed for the license which is shown in blue after this paragraph. Each license element will be shown on a separate line. If a license, or license element, is not valid for the current product then this will be shown in red.
HuEss-3.5-wcnp4-s4-BbftvACNMs-demmps-2010May31-f6d3931be921b6e1-(edwin@svi.nl)-c9cadd865d0d7b81
There were 2 possible reasons found why this license may not be valid to run this particular product.
Product name: HuEss
This license product name matches the currently running product.
Version: 3.5
The licensed verson is equal to the currently running product version.
Microscopes: wcnp4
'w' enables Widefield microscope functionality.
'c' enables Confocal microscope functionality.
'n' enables Nipkow disk microscope functionality.
'p' enables multi-photon microscope functionality.
'4' enables 4-Pi microscope functionality.
OK

FIGURE 2.4. The *Explain License* window lists all license details.

The license string as used by SVI has the same appearance on all supported platforms. For each product it is required to have a license string installed. Select a license string in the license window (HELP→LICENSE) and press the EXPLAIN LICENSE button. All details for the current license will be listed (Figure 2.4). If running into licensing problems this information can be used to analyze the problem.

Updating the Software

When the system is attached to the internet a pop-up window will appear when a newer version is available. The website can also be consulted for updates. Twice a year (April and October) new releases will become available. During and shortly after this period it is advisable to consult more frequently. Download the new version from the SVI website[2]. Proceed with the installation as explained above.

2. `http://www.svi.nl/`

Do not uninstall the old version as this will delete the license string. The newer version will by default automatically replace the older one. On Mac OS X please make a backup of the license string in a safe place before removing the previous installation.

Removing the Software

Removing the software will also cause the license string to be removed. If it is preferred to uninstall the current version prior to installing a newer one, take care to store the license string in a safe place. Table 2.2 on page 7 shows the uninstallation procedure for each platform.

TABLE 2.2. The uninstallation procedure per platform.

Platform	Procedure
Windows	Open the start menu and select: PROGRAMS→HUYGENS SUITE→UNINSTALL→REMOVE THE HUYGENS SUITE.
Linux	Open the package manager, search for *huygens* and uninstall it. This could also be handled with the command line; type `dpkg -r huygens` to install a Debian package or `rpm -e huygens` to install an RPM package.
Mac OS X	Drag the installed version to the waste basket.

System Requirements for Huygens Essential

Tables Table 2.3, Table 2.4, and Table 2.5 list the requirements for Windows, Mac OS X, and Linux.

TABLE 2.3. System requirements for Microsoft Windows.

Operating system	Huygens runs on Microsoft Windows 2003 Server, XP (32 and 64 bit), Vista (32 and 64 bit), and Windows 7 (32 and 64 bit)
Processor	AMD Athlon 64 or Intel Pentium 4 and higher.
Memory	2 Gb or more.
Graphics card	Any fairly modern card will do.

TABLE 2.4. System requirements for Mac OS X

Operating system	Huygens runs on Mac OS X Tiger (32 bit only) and (Snow)Leopard (32 and 64 bit)[a].
Processor	G5 PowerPC or Intel.
Memory	2 Gb or more.
Graphics card	Any fairly modern card will do.

a. OS X 10.5 or higher with X11 is required for full 64 bit capabilities.

TABLE 2.5. System requirements for Linux

Operating system	Most popular distributions like Ubuntu, RedHat, Fedora, and SuSE are supported (32 and 64 bit).
Processor	AMD Athlon 64 or Intel Pentium 4 and higher.

TABLE 2.5. System requirements for Linux

Memory	2 Gb or more.
Graphics card	Any fairly modern card will do.

Support on Installation

If any problem are encountered in installing the program or the licenses which could not be solved with the guidelines here included, please search the support Wiki[3] or contact SVI (See "Support and Contact Information" on page 108).

3. `http://www.svi.nl/FrontPage`

CHAPTER 3

The Image Restoration Process

The Processing Stages in the Wizard

Huygens Essential guides through the process of microscopic image deconvolution (also referred to as *restoration*) in several stages. Each stage is composed of one or more tasks. While proceeding, each stage is briefly described in the bottom-left *Task Info* window pane. The stage progress is indicated at the right side of the status bar. Additional information can be found in the online help (HELP→ONLINE HELP) as well as by clicking on the highlighted help questions.

The following steps and stages are to be followed:

- Loading an image.
- Stage *P*: the *preprocessing* stage. Here the possibility exists to load a microscopic parameter template, check the microscopic parameters, and crop the data.
- Stage 1: *parameter tuning*. This stage will be skipped if the preprocessing stage was aready passed. If the RESTART button is pressed in the last stage, then the wizard from stage 1 will be entered again.
- Stage 2: inspecting the image *histogram*.
- Stage 3: estimate the *background level*.
- Stage 4: the *deconvolution run*.
- Save the result.

The next sections will explain the stages in detail.

Loading an Image

Select FILE→OPEN to open the file dialog, browse to the directory where the images are stored, and select the image to be deconvolved, e.g. faba128.h5. A demo image (`faba128.h5`) is placed in the `Images` subdirectory of the installation path (see Table 2.1 on page 6).

Most file formats from microscope vendors are supported, but some of them require a special option in the license to be read. See the SVI support Wiki[1] for updated information.

1. `http://www.svi.nl/FrontPageFileFormats`

When the file is read successfully, either START DECONVOLUTION can be pressed to begin processing the image or the data can be converted using the tools in the TOOLS menu; some tools are described in the next subsections.

If a bead image was loaded, then one can also proceed selecting START PSF DISTILLER and proceed with generating a *point spread function* from measured beads (See Chapter 17 on page 105). A special license is needed in order to launch the PSF Distiller.

Additional images can be loaded for reference purposes (FILE→OPEN ADDITIONAL...), but only the one named *original* will be deconvolved during the guided restoration.

Converting a Dataset

Before pressing the START DECONVOLUTION button, a 3D stack can be converted into a 3D time series (TOOLS→CONVERT XYZ TO XYZT) or vice versa, or a 3D stack can be converted into a time series of 2D images (TOOLS→XYZ TO XYT) or vice versa.

Time Series

A time series is a sequence of images recorded along time at uniform time intervals. Every recorded image is a time frame. Huygens Essential is capable of automatic deconvolution of 2D-time or 3D-time data. There are some tools that are intended only for time series, as the confocal *bleaching corrector* or the *z-drift corrector*.

Verifying Microscopic Parameters

Next to the basic voxel data Huygens Essential also tries to read as much information as possible about the microscopic recording conditions. However, depending on the file type some information may be incomplete or missing. In this first stage all parameters relevant for deconvolution (Table 3.1 on page 10) are displayed and can be modified.

TABLE 3.1. Optical parameters explained.

Parameter	Explanation
Microscope type	Select from *widefield*, *confocal*, *spinning disk*, or *fourPi*.
Numerical aperture	The NA of the objective lens.
Objective quality	Select from perfect, poor, or something in between.
Coverslip position	The *position of the glass interface* between the immersion and embedding medium in µm, relative to the first slice of the stack.
Imaging direction	Select from *upward* or *downward*. Upward means that the objective lens is closest to the first slice in the stack.
Backprojected pinhole spacing	The distance (in µm) between the pinholes in the spinning disk as it appears in the *specimen plane*. This is the physical pinhole distance divided by the total magnification of the detection system.
Lens refractive index	The RI of the immersion medium for the objective lens.
Medium refractive index	The RI of the specimen embedding medium.
Backprojected pinhole radius	The radius (in nm) of the pinholes in the spinning disk as it appears in the *specimen plane*. This is the physical pinhole radius divided by the total magnification of the detection system.

TABLE 3.1. Optical parameters explained.

Parameter	Explanation
Excitation wavelength	The wavelength (in nm) of the excitation light (usually a laser line).
Emission wavelength	The wavelength (in nm) of the light emitted by the subject.
Excitation photon count	The number of photons used in *multi-photon* microscopy.
Excitation fill factor	The width of the beam relative to the aperture. The default for this value is 2, meaning that the aperture has a diameter of 2σ, where σ is the standard deviation of the Gaussian distribution in the beam.

If values are displayed in a *red* background, they are highly suspicious. An *orange* background indicates a non-optimal situation (See Figure 3.1). Oversampling is also indi-

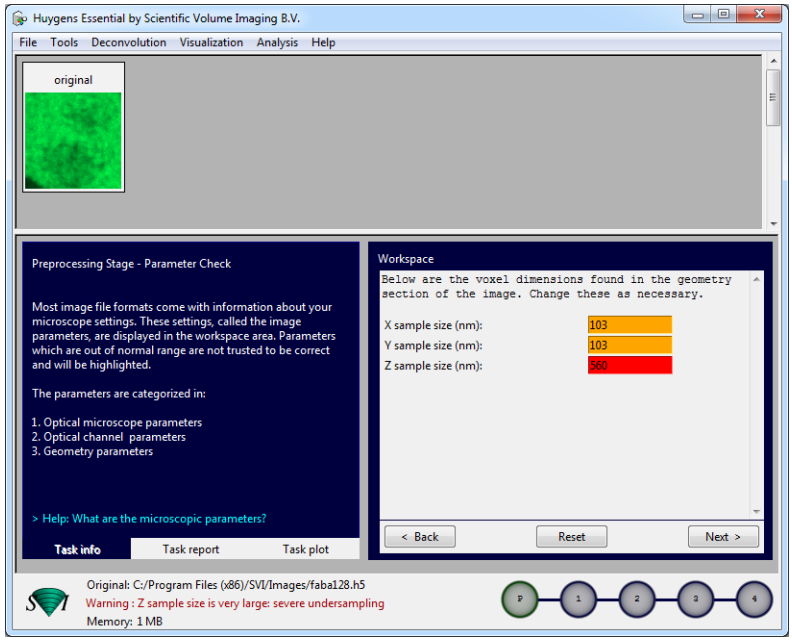
Huygens Essential by Scientific Volume Imaging B.V.
File Tools Deconvolution Visualization Analysis Help
original
Preprocessing Stage - Parameter Check
Most image file formats come with information about your microscope settings. These settings, called the image parameters, are displayed in the workspace area. Parameters which are out of normal range are not trusted to be correct and will be highlighted.
The parameters are categorized in:
1. Optical microscope parameters
2. Optical channel parameters
3. Geometry parameters
> Help: What are the microscopic parameters?
Task info
Task report
Task plot
Workspace
Below are the voxel dimensions found in the geometry section of the image. Change these as necessary.
X sample size (nm): 103
Y sample size (nm): 103
Z sample size (nm): 560
< Back
Reset
Next >
Original: C:/Program Files (x86)/SVI/Images/faba128.h5
Warning : Z sample size is very large: severe undersampling
Memory: 1 MB

FIGURE 3.1. Parameter check stage: Sampling. Red coloring indicates a suspicious value, and orange a non-optimal value.

cated with a *cyan* background, that becomes *violet* when it is very severe.

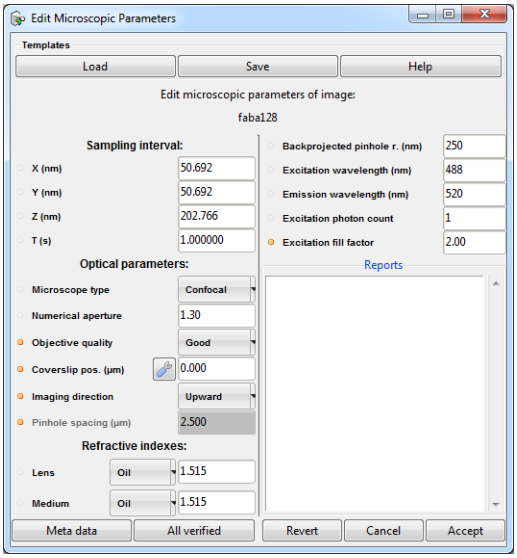

FIGURE 3.2. The microscopic parameter editor. This window can be opened by right-clicking on the thumbnail and selecting EDIT PARAMETERS.

The image parameters can be checked and corrected, not only at this deconvolution stage, but also at any time by right-clicking on the image thumbnail and selecting SHOW PARAMETERS or EDIT PARAMETERS. The parameter editor is show in Figure 3.2.

Microscopic Template Files

Once the proper parameters have been set and verified, they can be saved to a Huygens template file (`.hgst`). These templates can be applied at the start of the wizard, hence the user can skip the parameters verification stage, provided that an image is to be restored with the same optical properties as the ones which were recorded on the template.

The LOAD MICROSCOPIC TEMPLATE button will allow the selection of a template from a list of saved template files which reside both in the common templates directory and in the user's personal template directory. The Huygens common templates directory is named `Templates`, and resides in the Huygens installation directory (See Table 2.1 on page 6). The user's personal templates directory is called `SVI/Templates` and be found in the user's home directory[2].

The Intelligent Cropper

The time needed to deconvolve an image increases more than proportional with its volume. Therefore, deconvolution can be accelerated considerably by cropping the image. Huygens Essential is equipped with an intelligent cropper which automatically surveys the image to find a reasonable proposal for the crop region (See Figure 3.3). In computing this initial proposal the microscopical parameters are taken into account, making sure that cropping will not have a negative impact on the deconvolution result. Because the survey depends on accurate microscopical parameters it is recommended to use the cropper as final step in the preprocessing stage (press YES when the wizard asks to launch the cropper), but it can be launch from outside the wizard through the menu TOOLS→CROP.

Cropping in *X*, *Y*, and *Z*.

The borders of the proposed cropping region are indicated by a red contour. The initial position is computed from the image content and the microscopic parameters at launch time of the cropper.

2. The user home directories are usually located in `C:\Users` on Windows 7 and Vista and in `C:\Documents and Settings` on Windows XP and lower. On Mac OS X they are usually in `/Users` and on Linux in `/home`.

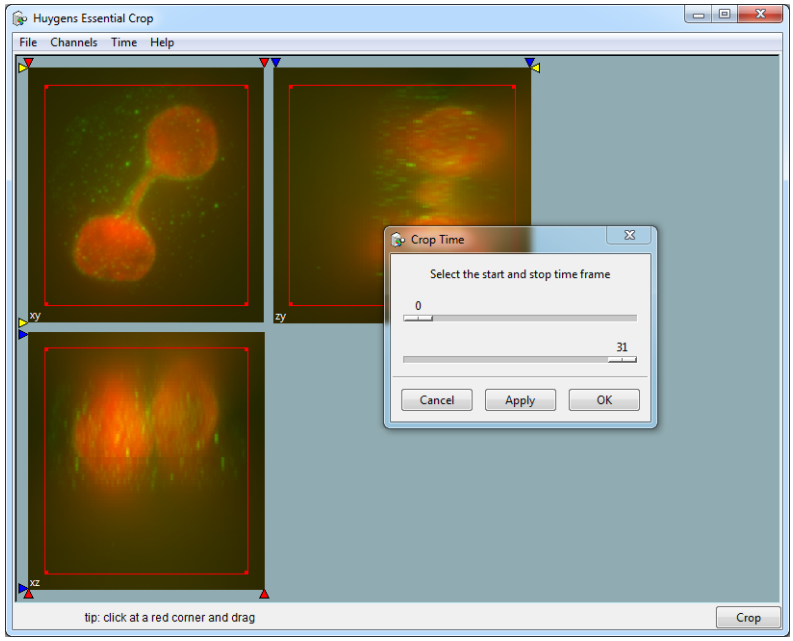

FIGURE 3.3. The crop tool in Huygens Essential.

The three views shown are *maximum intensity projections* (MIPs) along the main axes. By default the entire volume (including all time frames) is projected. The red, yellow, and blue triangles can be dragged to restrict the projected volume.

The cropper allows manual adjustment of the proposed crop region. To adjust the crop region put the cursor inside the red boundary, press the left mouse button and drag the contour to the preferred position. Accept the new borders by pressing the CROP button. Do not crop the object too tightly, because that would remove blur information relevant for deconvolution.

Cropping in Time

The number of frames in a time series can be reduced by selecting TIME→SELECT FRAMES... from the cropper menu.

Removing Channels

The number of channels in a multi-channel image can be reduced by selecting CHANNELS→SELECT CHANNELS... from the cropper menu.

The Image Histogram

The histogram is an important statistical tool for inspecting the image. It is included to be able to spot problems that might have occurred *during the recording*. It has no image manipulation options as such, it just may be preventing from future recording problems.

The histogram shows the number of pixels as a function of the intensity (gray value) or groups of intensities. If the image is an 8 bit image (gray values in the range 0-255) the *x*-axis is the gray value and the *y*-axis is the number of pixels in the image with that gray value. If the image is more than 8 bit, then gray values are collected to form a *bin*. For

example, gray values in the range 0-9 are collected in bin 0, values in the range 10-19 in bin 1, etc. The histogram plot now shows the number of pixels in every bin.

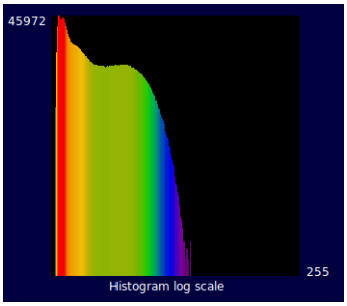

FIGURE 3.4. The image histogram. The vertical mapping mode can be selected from linear, logarithmic or sigmoid.

The histogram in Figure 3.4 shows that the intensity distribution in the demo image is of reasonable quality. The narrow peak shown at the left represents the background pixels, all with similar values. The height of the peak represents the amount of background pixels (note that the vertical axes uses *logarithmic* scaling). Because in this particular image there are many voxels with a value in a narrow range around the background the peak is higher than the other.

In this case there is also a small black gap at the left of the histogram. This indicates an electronic offset, often referred to as *black level*, in the signal recording chain of the microscope.

If a peak is visible at the extreme right hand side of the histogram it indicates *saturation* or *clipping*. Clipping is caused by intensities above the maximum digital value available in the microscope. Usually, all values above the maximum value are replaced by the maximum value. On rare occasions they are replaced by zeros. Clipping will have a negative effect on the results of deconvolution, especially with widefield images.

The histogram stage is included for examining purpose only. It does not affect the deconvolution process that follows.

Estimating the Average Background

In this stage the *average* background in a volume image is estimated. The average background corresponds with the noise-free equivalent of the background in the measured (noisy) image. It is important for the search strategy that the microscopic parameters of the image are correct, in especially the sampling distance and the microscope type.

The following search strategies are available:

- **Lowest value** (default): The image is searched for a 3D region with the *lowest average value*. The axial size of the region is about 0.3 µm; the lateral size is controlled by the radius parameter which is by default set to 0.5 µm.
- **In/near object**: The *neighborhood around the voxel with the highest value* is searched for a planar region with the lowest average value. The size of the region is controlled by the radius parameter.
- **Widefield**: First the image is searched for a 3D region with the lowest values to ensure that the region with the least amount of blur contributions is found. Subsequently the background is determined by searching this region for the planar region with radius *r* that has the lowest value.

Press the ESTIMATE button in the wizard to continue. If the estimated value should be checked, then open the image in the *Twin Slicer* and hover over a background area; the intensity values are displayed at the top. The value could now be adapted either by altering the value in the *Estimated background* field or in the *Relative background* field. Setting the latter to -10, for example, lowers the estimated background by 10%. If done press ACCEPT to proceed to the deconvolution stage.

The Deconvolution Stage

Huygens Essential uses *Classical Maximum Likelihood Estimation* for the deconvolution process[3]. This method is extremely versatile; applicable for all types of data sets. The following parameters to this algorithm can be set:

1. **Number of iterations**. MLE is an iterative process that never stops if no *stopping criterion* is given. This stopping criterion can simply be the maximum number of iterations. This value depends on the desired final quality of the image. For an initial run the value can be left at its default. To achieve the best result this value can be increased to e.g. 100. Another stopping criterion is the *Quality threshold* of the process (See Item 3).
2. **Signal to noise ratio**. The SNR is a parameter than controls the sharpness of the restoration result. Using a too large SNR value might be risky when restoring noisy originals, because the noise could just being enhanced. A noise-free widefield image usually has SNR values higher than 50. A noisy confocal image can have values lower than 20.
3. **Quality threshold**. Beyond a certain amount of iterations, typically below 100, the difference between successive iterations becomes insignificant and progress grinds to a halt. Therefore it is a good idea to monitor progress with a quality measure, and to stop iterations when the change in quality drops below a threshold. At a high setting of this quality threshold, e.g. 0.1, the quality difference between subsequent iterations may drop below the threshold before the indicated maximum number of iterations has been completed. The smaller the threshold the larger the number of iterations which are completed and the higher the quality of restoration. Still, the extra quality gain becomes very small at higher iteration counts.
4. **Iteration mode**. In *optimized* mode (highly recommended) the iteration steps are bigger than in *classical* mode. The advantage of classical mode is that the direction of its smaller steps is sure to be in the right direction; this is not always the case in optimized mode. Fortunately, the algorithm detects if the optimized mode hits upon a sub optimal result. If so, it switches back to the classical mode to search for the optimum.
5. **Bleaching correction**. If this option is set to *if possible*, then the data is inspected for bleaching. 3D stacks and time series of widefield images will always be corrected. Confocal images can only be corrected if they are part of a time series, and when the bleaching over time shows exponential behavior.
6. **Brick layout**. When this option is set to *auto*, then Huygens Essential splits the image in bricks in two situations:
   a. The system's *memory* is not sufficiently large to allow an image to be deconvolved as a whole.
   b. *Spherical aberration* is present, for which the point spread function needs to be adapted to the depth.

Press DECONVOLVE to start the restoration process (See Figure 3.5). Pressing STOP halts the iterations and retrieves the result from the previous iteration. If the first iteration is not yet complete a empty image will result.

3. Huygens Professional also has *Quick-MLE-time*, *Quick-Tikhonov-Miller*, and *Iterative Constrained Tikhonov-Miller* algorithms available.

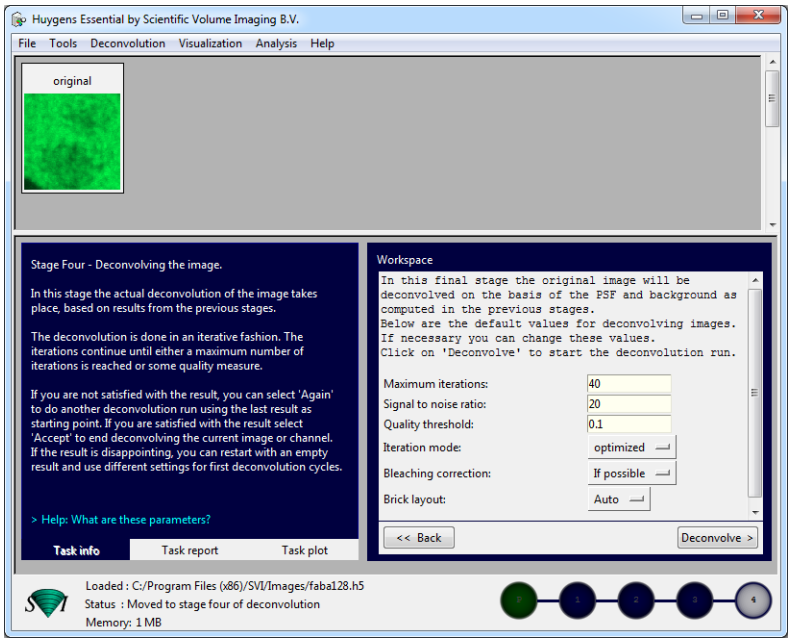
Huygens Essential by Scientific Volume Imaging B.V.
File Tools Deconvolution Visualization Analysis Help
original
Stage Four - Deconvolving the image.
In this stage the actual deconvolution of the image takes place, based on results from the previous stages.
The deconvolution is done in an iterative fashion. The iterations continue until either a maximum number of iterations is reached or some quality measure.
If you are not satisfied with the result, you can select 'Again' to do another deconvolution run using the last result as starting point. If you are satisfied with the result select 'Accept' to end deconvolving the current image or channel. If the result is disappointing, you can restart with an empty result and use different settings for first deconvolution cycles.
> Help: What are these parameters?
Task info
Task report
Task plot
Workspace
In this final stage the original image will be deconvolved on the basis of the PSF and background as computed in the previous stages.
Below are the default values for deconvolving images.
If necessary you can change these values.
Click on 'Deconvolve' to start the deconvolution run.
Maximum iterations: 40
Signal to noise ratio: 20
Quality threshold: 0.1
Iteration mode: optimized
Bleaching correction: If possible
Brick layout: Auto
<< Back
Deconvolve >
Loaded : C:/Program Files (x86)/SVI/Images/faba128.h5
Status : Moved to stage four of deconvolution
Memory: 1 MB

FIGURE 3.5. The Deconvolution Stage in the wizard.

Finishing or Restarting a Deconvolution Run

When a deconvolution run is finished use the *Twin Slicer* to inspect the result in detail. Depending on the outcome of it there can be selected to RESTART, RESUME or ACCEPT the restoration:

- **Restart** discards the present result, and returns to the very first stage where the microscopic parameters can be entered. Now the process can be restarted with different microscopic and/or deconvolution parameters.
- **Resume** keeps the result and returns to the stage where the deconvolution parameters can be entered. The software will ask to continue at the left off, or to start from the raw image again. A new result will be generated to compare with the previous one. This can be repeated several times.
- **Accept** proceeds to the final stage or, if the data was multi-channel, to the next channel. If several results are generated by resuming the deconvolution there will be asked to select the best result as the final one, that will be renamed to *deconvolved*. The other results will remain as well in case it is desired to save them.

Multi-channel Images

Multi channel images can be deconvolved in a semi automatic fashion, to give the opportunity to fine tune the results obtained with each individual channel. After the preprocessing stage the multi channel image is split into single channel images named *channel-0*, *channel-1*, etc. The first of these is automatically selected for deconvolution.

The procedure to deconvolve a channel in a multi channel data set is exactly the same as for a single channel image. Therefore multiple reruns on the channel can be done at at hand, just as with single channel data. When everything is done press ACCEPT in the last stage. This will cause the next channel to be selected for restoration. Proceed as usual with that channel and the remaining channels. If it is not needed to process all the channels in an image one or more channels may be skipped.

When the last channel has been processed, the wizard allows to select the results which yshould be combine into the final deconvolved multi channel image. This means that up to this point it is still possible to decide which of the results to combine, even in what order. Once ACCEPT is pressed a multi channel image named *deconvolved* is created.

Z-drift Correcting for Time Series

For 3D time series the wizard shown an additional stage to enable correction for movement in the *z* direction (axial) that could have been occurred for instance by thermal drift of the microscope table. In case of a multi channel image, the corrector can survey *All channels* and determine the mean *z* position of the channels, or it can take *One channel* as set by the *Reference channel* parameter.

After determining the *z* positions per frame, the *z* positions (not the image) can be filtered using a *median*, *Gaussian* or *Kuwahara* filter of variable width. When the drift is gradual, a Gaussian filter is probably best. In case of a drift with sudden reversals or outliers a median filter is best. In case the *z* positions show sudden jumps, we recommend the Kuwahara filter.

Saving the Result

After each deconvolution run the result can be saved. Select the image to be saved and select FILE→SAVE 'IMAGENAME' AS... in the menu bar. The *HDF5* file format preserves all microscopic parameters and applies a lossless compression.

Select DECONVOLUTION→SAVE TASK REPORT to store the information as displayed in the Task report tab.

Using a Measured PSF

Measured PSF's improve deconvolution results and may also serve as a quality test for the microscope. If a PSF was loaded (FILE→OPEN PSF), then Huygens Essential will automatically use it. If the measured PSF contains less channels than the image, a theoretical PSF will be generated for the channels where there is no PSF available. See “The PSF Distiller” on page 105 and “The Point Spread Function” on page 102 for more information.

CHAPTER 4

The Twin Slicer

The Twin Slicer allows to synchronize views of two images, measure distances, plot line profiles, etc. In *basic mode*, which is also available without a license, image comparison is intuitive and easy, while the *advanced mode* gives the user the freedom to rotate the cutting plane to any arbitrary orientation, link (synchronize) or unlink viewing parameters between the two images, and more.

To launch the Huygens Twin Slicer, select an image and select VISUALIZATION→TWIN SLICER from the main menu. To view another image in an existing slicer, click the image name in the drop-down menu above the left or right view port (See Figure 4.1).

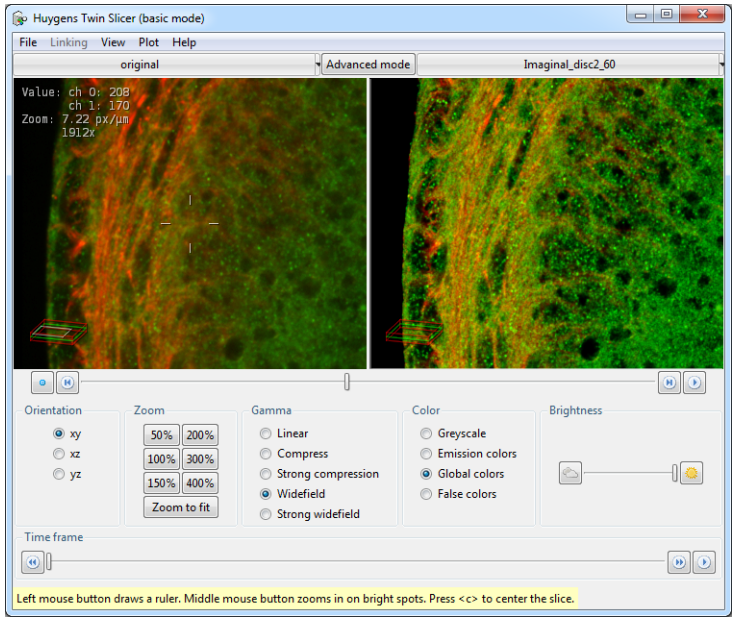

FIGURE 4.1. The Twin Slicer in *basic mode*, showing an original and deconvolved image side-by-side.

The View Menu

Use the VIEW menu to show or hide image properties and guides. These are listed in Table 4.1:

TABLE 4.1. The options in the Twin Slicer's VIEW menu.

Option	Description
POINTER COORDINATES	Display the position of the mouse pointer in µm or in voxel coordinates.
TIME	Display the time for the current slice in seconds or frame numbers.
INTENSITY VALUES	Display the intensity values for all channels on the current pointer location.
ZOOM	Display the zoom value in screen-pixels per micron. A magnification factor is displayed as well; using the pixel density for the monitor, this value gives an estimation for the absolute magnification.
ROTATION ANGLES	Display the tilt and twist angles in degrees.
DROP SHADOWS	Enhance the contrast for the overlayed lines and text by showing drop shadows.
SLICE BOUNDARIES	Draws the slice boundaries for the left image in the right one and vice versa. This is helpful when both slicers are used.
WIREFRAME BOX	Show or hide the wireframe box, which gives visual feedback on the position and orientation of the cutting plane (green), and the displayed slice (gray) in the data volume (red).
SVI LOGO	Show or hide the SVI logo in the lower right of the view port.

Panning

Click and hold the *right mouse button* on the slice to move it around. Clicking the *center* button ( ) or pressing the 'c' key centers the slice.

Slicing

Drag the *slider below the view ports* to move the cutting plane back and forth. This can also be achieved using the buttons adjacent to the slider (and), the up/down arrow keys on the keyboard, or by placing the mouse pointer over the slider and using the scroll wheel. The play button ( ) moves the cutting plane through the data volume. The pointer coordinates can be displayed through the VIEW menu. Note that it is possible to move the cutting plane out of the volume. Pressing the *center* button ( ) or pressing the 'c' key centers the plane again.

Using the Slicer in Basic Mode

The button centered at the top of the window enables switching between *basic* and *advanced* mode. In basic mode, all controls are visible in the panels below the view ports (See Figure 4.1). In contrast to the advanced mode, which allows independent control of the left and right slicer (See "Using the Slicer in Advanced Mode" on page 22), the basic mode shows a single set of controls that apply to both slicers.

Changing Time Frames

Drag the slider in the lower *Time frame* panel to change the time frame or press the play button ( ) to animate the time series. The time frame can be displayed through the VIEW menu.

Orientation

Make a selection in the most left *Orientation* panel to change the plane used to display the image.

Zooming

Click the buttons in the *Zoom* panel or use the scroll wheel to zoom in or out on the location of the mouse pointer.

Changing Display Colors

Click an option in the *Color* panel to select a color scheme:

- **Greyscale**: the image is displayed in gray tints. For single-channel images, this gives a higher contrast than the emission or global colors.
- **Emission colors**: if the emission wavelengths are set correctly, this gives the most intuitive view.
- **False colors**: a false color is given to each intensity value. This view gives a high contrast and makes it easy to spot areas of homogeneous intensity.
- **Global colors**: the colors as defined in the *global color scheme*. The global color scheme applies to all visualization tools and can be modified via the Huygens Essential main menu: TOOLS→PREFERENCES...→EDIT GLOBAL COLORS.
- **Custom**

Tuning the Brightness and Contrast

The brightness can be adjusted in the most right *Brightness* panel using the buttons (and), dragging the slider, or putting the mouse pointer over the slider and using the scroll wheel. The *Gamma* panel provides a linear and some nonlinear ways of mapping data values to pixel intensities. These are:

- **Linear** (default): pixel values are mapped to screen buffer color intensities in a linear fashion. Note that the actual translation of the screen buffer values to the actual brightness of a screen pixel is usually quite nonlinear.
- **Compress**: where an image contains a few very bright spots and some larger darker structures using linear mode will result in poor visibility of the darker structures. Restoration of such images is likely to further increase the dynamic range resulting in the large structures becoming even dimmer. In such cases use the compress display mode to increase the contrast of the low valued regions and reduce the contrast of the high-valued regions. Another way to improve the visibility of dark structures is the usage of *false colors* (See "Changing Display Colors" on page 21).
- **Widefield**: in restoring widefield images it sometimes happens that blur removal is not perfect, for instance when one is forced to use a theoretical point spread function in sub optimal optical conditions. In such cases the visibility of blur remnants can be effectively suppressed.

Automatic Panning, Slicing and Zooming

When the *middle mouse button is clicked*, the Twin Slicer will automatically center and zoom in on the brightest spot in a 3D neighborhood around the cursor.

Using the Slicer in Advanced Mode

The button centered at the top of the window offers switching between *basic* and *advanced* mode. The advanced mode allows independent control of the left and right slicer. In this mode, all controls are available in twofold and accessible through the tabs in the bottom of the window.

Changing Time Frames

Drag the slider in the *Time frame* tab to change the time frame or press the play button ( ) to animate the time series. The time frame can be displayed through the VIEW menu.

Zooming

Use the scroll wheel to zoom in or out on the location of the mouse pointer, or access the *Zoom* tab. The four buttons in this tab respectively *zoom out* ( ), *zoom in* ( ), *zoom 1:1* ( ) (the x-sample distance matches 1 pixel), and *view all* ( ).

Rotation

The three radio buttons in the *Rotate* tab can be used to switch between axial (*xy*), frontal (*xz*), and transverse (*yz*) orientations. The *Twist* slider rotates the cutting plane around a *z*-axis, while the *Tilt* button rotates the cutting plane around an axis in the *xy* plane. The tilt and twist angles can be displayed through the VIEW menu. Note that the wireframe box in the bottom left of each view port gives visual feedback about the position and orientation of the slice.

Changing Display Colors

Click the *Colors* tab key to view the color settings panel. The *Active channels* drop down menu can be used to enable or disable channels.

In addition to the *color schemes* that are available in basic mode (“Changing Display Colors” on page 21), the advanced mode allows the use of *custom colors*. Use the color picker ( ) to manually select a color for each channel.

Tuning the Brightness and Contrast

The brightness and contrast controls are accessible in the *Contrast* panel. The brightness can be changed per channel, or for all channels at once. The *Gamma* drop down menu provides a linear and some non-linear ways of mapping data values to pixel intensities (See “Tuning the Brightness and Contrast” on page 21 for an overview).

If the *Link channels* box is checked, this means that the way of mapping data values to pixel intensities is the same for all channels; if not, the range is automatically adjusted for to minimum and maximum in each channel.

Linking Controls

The LINKING menu can be used to change the way in which both slicers communicate. The options in this menu are listed in Table 4.2. Note that settings get synchronized once the controls are being used.

TABLE 4.2. The options in the Twin Slicer's LINKING menu (accessible in advanced mode).

Option	Description
ADVANCED LINKING	Enables the user to change the linking of the slice position, panning, and rotation. Doing so may lead to complex situations regarding orientation.
POINTER LOCATION	Shows the position of the mouse pointer in the other slicer.
SLICE POSITION	Makes sure that the cutting plane for the right slicer crosses the center of the left slice, and vice versa.
TIME FRAME	Synchronize the time.
ZOOM LEVEL	Synchronize the level of magnification.
PANNING	This does not affect position of the cutting plane, but it shifts the right slice such that the projection of the center of the left slice is in the center of the right slice, and vice versa.
ROTATION	Makes sure that the rotation angles for both cutting planes are the same.
ACTIVE CHANNELS	The left and right slicer will have the same channels enabled and disabled.
COLOR SCHEME	Makes sure that the left and right slicer use the same colors scheme.
CUSTOM COLORS	Use the same custom color scheme for both slicers.
BRIGHTNESS	Synchronize the brightness.
GAMMA	Synchronize the gamma setting.

Some useful ways of linking the controls are:

- **Comparison mode**: to configure the Huygens Twin Slicer to compare two images, e.g. original and deconvolved, it is best to link all orientation parameters, i.e. slice position, time frame, zoom level, panning and rotation. This ensures that there is always looked at the same piece of data.
- **Orthogonal mode**: to view a part of an image in two orthogonal directions, for instance axial (*xy*) and frontal (*xz*), do the following:
  - Select the same image for both the left and right slicer.
  - Tick ADVANCED LINKING and link the slice position, time frame, zoom level, and panning. Unlink the rotation.
  - Select the *Rotate* tab at the bottom of the window and select the *xz* and *xz* orientation.

  Now it is possible to zoom, pan, and slice while the centers of the left and right slice are always aligned. Note that when the cutting planes are not the same, the projected mouse pointer will show a distance (in μm) besides it. If this number is positive, it means that real pointer is more towards the observer (in front of the screen).
- **Overview mode**: An easy overview mode can be configured as follows:
  - Select the same image for both the left and right slicer.

- Tick ADVANCED LINKING and link the slice position, time frame, and rotation. Unlink the zoom level and panning.
- Drag the sash to the left to make the left slicer a bit smaller.
- Select the *Zoom* tab at the bottom and click the *view all* button ( ).

Now the right slicer can be used to zoom in on the data, while the left slicer shows the position in the image (See Figure 4.2).

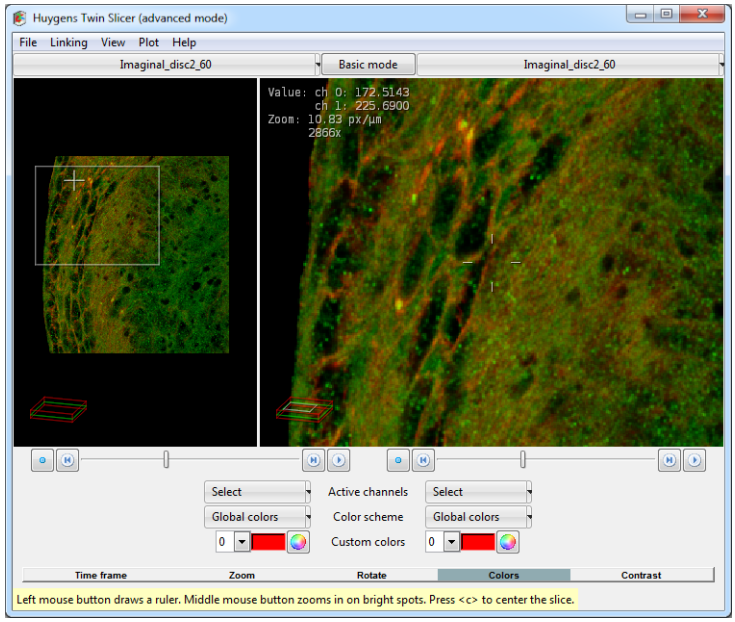

FIGURE 4.2. The Twin Slicer in *advanced mode*, with all controls but zoom and panning linked.

Measurement

Markers

Double click in one of the images to place a marker at the position of the mouse pointer. As configured in the VIEW menu, the marker shows the coordinates and intensity values besides it. To remove the marker, click it and press the `Delete` key.

Rulers

To overlay a ruler on the image, *hold the left mouse button and drag*. The length of the line in µm is displayed beside it. Click and drag the end points of the ruler to make adjustments. Press and hold the `Ctrl` key while dragging an end point to change length without changing direction. Click and drag the middle of the ruler to move it in its entirety, without changing length or direction. Press and hold the `Ctrl` key while dragging the ruler to move it perpendicular to its direction. To remove the ruler, click it and press the `Delete` key.

Intensity Profiles

When a ruler in the left slicer is selected, the right slicer will be replaced by a plot window and vice versa. See the online SVI support Wiki[1] for more information on the data plotter.

Select PLOT→PLOT BOTH SLICERS from the menu to show the intensity profiles for both the left and right image in the same plot. The graphs for the left slicer will have solid lines, while the graphs for the right one are dashed (See Figure 4.3).

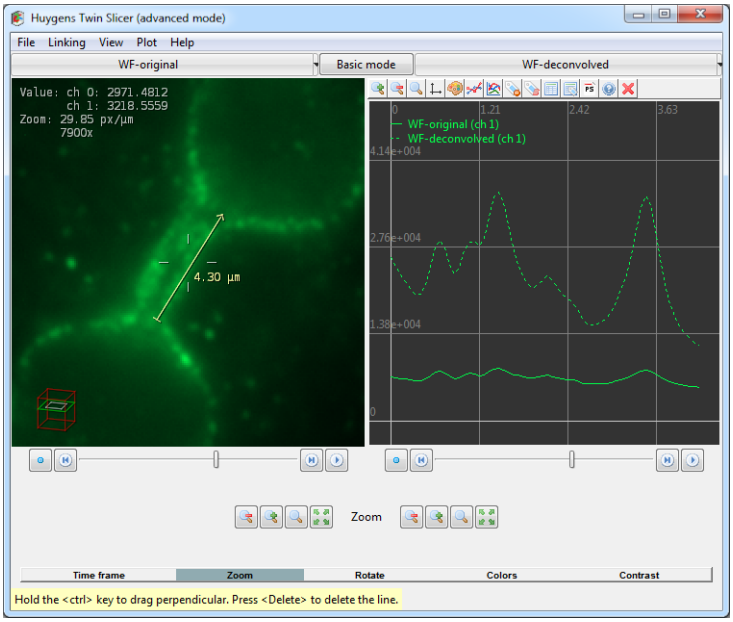

FIGURE 4.3. Measuring the intensity profile along a line. The plot can be configured such that it shows the profile of both images (left solid, right dashed).

1. http://www.svi.nl/FrontPageDataPlotter

CHAPTER 5

The Orthogonal Slicer

The Huygens Orthogonal Slicer, shown in Figure 5.1, is designed to show the same point in 3D space from three orthogonal directions;

- axial or *xy* (top left);
- frontal or *xz* (bottom left);
- transverse or *yz* (bottom right).

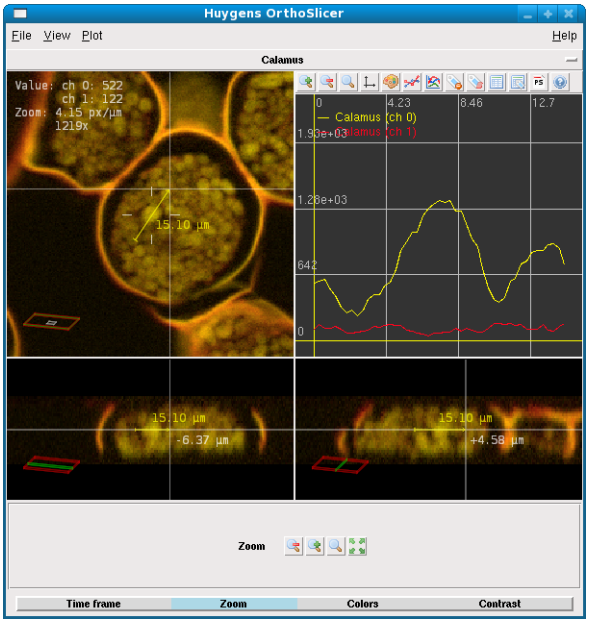

FIGURE 5.1. The Huygens Orthogonal Slicer

If you move one of the slices, the others will follow to make sure that the center of each of the slices intersects in the same point in space. This behaviour makes the Ortho Slicer a useful tool to study small objects in 3D.

The Crosshair Cursor

The position of your mouse is projected as a cross-hairs pointer on all slices. The value besides the center of the cross-hairs gives the distance of the mouse position to this projection. If this number is positive, it means that real pointer is more towards you (in front of your screen).

Visualization parameters

Changing the visualization parameters in the Orthogonal Slicer is similar to the Huygens Twin Slicer (on page 19). There are tools to;

- change time frames
- zoom in, out, fit or zoom 1:1
- change display colors
- tune the brightness and contrast

Panning can be achieved by right-clicking and dragging an image. To center the slice, press the Center button (the blue dot) at the lower left of the image or press 'c'.

Measurements

To overlay a ruler on the image, hold the left mouse button and drag. The length of the line in microns is displayed beside it. Left-click and drag the end points of the ruler to make adjustments. Note that the other orthogonal directions show a projection of this ruler. Press and hold `Ctrl` while dragging an end point to change length without changing direction. Left-click and drag the middle of the ruler to move it in its entirety, without changing length or direction. Press and hold the `Ctrl` key while dragging the ruler to move it perpendicular to its direction. To remove the ruler, left-click somewhere else on the image.

When a ruler is drawn, the help pane will be replaced by a plot that shows the intensity profile along it. See the SVI wiki for more information about the data plotter's capabilities.[1]

Auto-Zoom

When you click the middle mouse button, the Orthogonal Slicer will automatically center and zoom in on the brightest spot in a 3D neighborhood around the mouse pointer.

Display Options

The VIEW menu allows you to show or hide information and guides within the image overlay, including pointer coordinates, time, intensity, zoom, rotation, and the wireframe.

The 'Global value range' option in the PLOT menu uses the maximum and minimum value of the image(s) to determine the visible range of the plot, otherwise it uses the maximum and minimum values of the plot data.

1. `http://support.svi.nl/DataPlotter`

CHAPTER 6

The MIP Renderer

The *Maximum Intensity Projection* (MIP) Renderer enables the possibility to obtain an orthogonal projection of 3D data from any given viewpoint.

The renderer projects the voxels with maximum intensity that fall in the way of parallel rays traced from inside the image volume to the screen (See Figure 6.1). Notice that this implies that two MIP renderings from opposite viewpoints show symmetrical images.

To start the MIP Renderer, right-click on a thumbnail and select VIEW→MIP RENDERER from the pop-up menu.

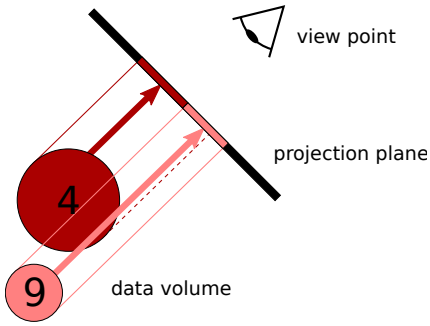

FIGURE 6.1. A schematic overview of MIP rendering. The maximum intensities on rays perpendicular to the screen are projected.

Basic Usage

Orientation and Zoom

Adjust the viewpoint by moving the *Tilt* and *Twist* sliders (See Figure 6.2), or by dragging the mouse pointer on the large view. The magnification can be adjusted using the *Zoom* slider or by using the scroll wheel. Use right mouse button to *pan* the center of the projection.

Note that the thumbnail preview (the top right) reflects changes in the configuration instantly, while the large view should be updated manually. To update the large view, press the *fast mode* button ( ) or the *high quality* button ( ).

Threshold

The *Soft threshold* slider in the *Channel parameters* panel at the right affects the threshold level. The application of a threshold is a preprocessing step that *reduces the background* in the image, i.e. voxels with intensity values below the threshold value become transparent. Contrary to a standard threshold, which is 'all or nothing' (values above the

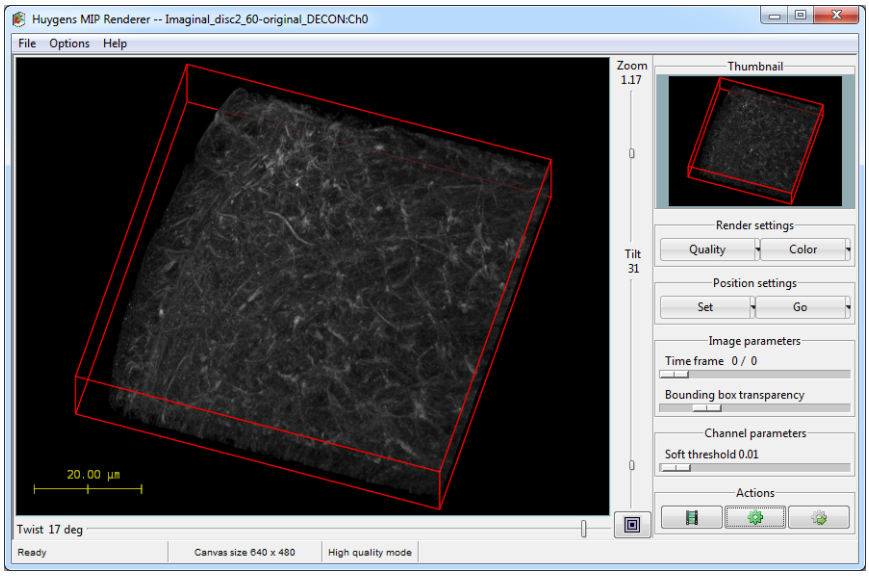

FIGURE 6.2. The MIP Renderer window.

threshold are kept, values below it are deleted), the soft threshold function handles images in a different way. It makes a *smooth transition* between the original an the deleted value.

Saving Scenes

Choose FILE→SAVE SCENE... to save the rendered scene as a Tiff file.

Advanced Usage

Render Options

Table 6.1 gives an overview of the different render options that are available through the

TABLE 6.1. Render options for the MIP Renderer.

Option	Description
ANIMATION FRAME COUNT	Set the number of frames that will be rendered in a movie. 180 frames with a frame rate of 24 fps result in a movie with a duration of 7.5 seconds.
ANIMATION FRAME RATE	Adjust the frame rate; a rate of 24 frames per second is fine for smooth movies.
RENDER SIZE	Adjust the size of the rendered image. When the render size exceeds the display area, then use the *middle mouse button* to pick up and move the rendered image.
RENDER QUALITY	Set the default quality (FAST or HIGH QUALITY). This setting will be used for rendering animations.
COLOR MODE	Choose between GREY, EMISSION COLORS, GLOBAL PALETTE (See "Adjusting the Global Color Scheme" on page 104), or FALSE COLOR.
BOUNDING BOX	Enable or disable the bounding box, or adjust the line color.
SHOW SCALE BAR	Enable or disable the scale bar.
SOFT THRESHOLD MODE	Adjust the smoothness of the soft threshold (See "Threshold" on page 29).
SHOW SVI LOGO	Hide or show the SVI logo at the bottom right.
CENTER SCENE	Undo both the panning of the projection center (right mouse button) and the rendered image itself (middle mouse button).

OPTIONS menu. The ANIMATION FRAME COUNT, ANIMATION FRAME RATE and RENDER QUALITY apply to the rendering of simple movies as explained in the next section.

Templates

All scene settings, i.e. both the render options and all parameters, can be exported to a template file via FILE→SAVE SCENE TEMPLATE. The template files have the extension `.hgsv` and they can be applied to any image that is loaded in the MIP Renderer.

Simple Animations

The Huygens *Movie Maker* (See "The Movie Maker" on page 43) allows to create easily sophisticated animations using the MIP, SFP, and Surface Renderer.

Without the Movie Maker the MIP Renderer has the option to make simple animations of the image, changing the view point in different frames. Set the render parameters for the first frame and click SET→HOME in the *Position settings* panel at the right. Now adjust the viewpoint for the final frame, and click SET→END. Also the *frame count*, *frame rate*, or other render options in the OPTIONS menu may be adjusted. Finally press the *animate* button (), and select a directory to save the AVI movie or the TIFF frames to.

The exported AVI files use the MJPEG[1] codec and can be loaded in most movie players, including Windows Movie Player and Apple Quicktime. TIFF frames are useful to combine multiple animations or edit the movie in e.g. Windows Movie Maker.

1. http://en.wikipedia.org/wiki/Mjpeg

CHAPTER 7

The SFP Renderer

The SFP Renderer generates *realistic 3D scenes*, based on taking the 3D microscopy image as a distribution of fluorescent material. The computational work is done by the *Simulated Fluorescence Process* (SFP) algorithm[1], simulating what happens if the material is excited and how the subsequently emitted light travels to the observer (See Figure 7.1). The unique properties of this algorithm enable it to create depth cue rich images from unprocessed data. Because it does not rely on boundaries or sharp gradients, it is eminently suited to render 3D microscopic data sets. Since the SFP algorithm is based on *ray tracing* that runs efficiently on multi-core computers, it does not require a special graphics card.

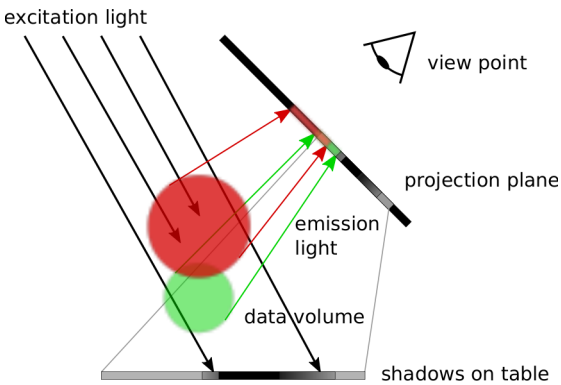

FIGURE 7.1. In the SFP renderer excitation and subsequent emission of light of fluorescent materials is simulated. Each subsequent voxel in the light beam is affected by shadowing from its predecessors. The transparency of the object for the emission light controls to what extent the viewer can peer inside the object.

To start the SFP Renderer, right-click on a thumbnail and select VIEW→SFP RENDERER from the pop-up menu, or choose VISUALIZATION→SFP RENDERER from the main menu.

Basic Usage

Orientation and Zoom

Adjust the viewpoint by moving the *Tilt* and *Twist* sliders (See Figure 7.2) or by dragging the mouse pointer on the large view. The magnification can be adjusted using the *Zoom* slider or by using the scroll wheel.

1. http://www.svi.nl/FrontPageSFP

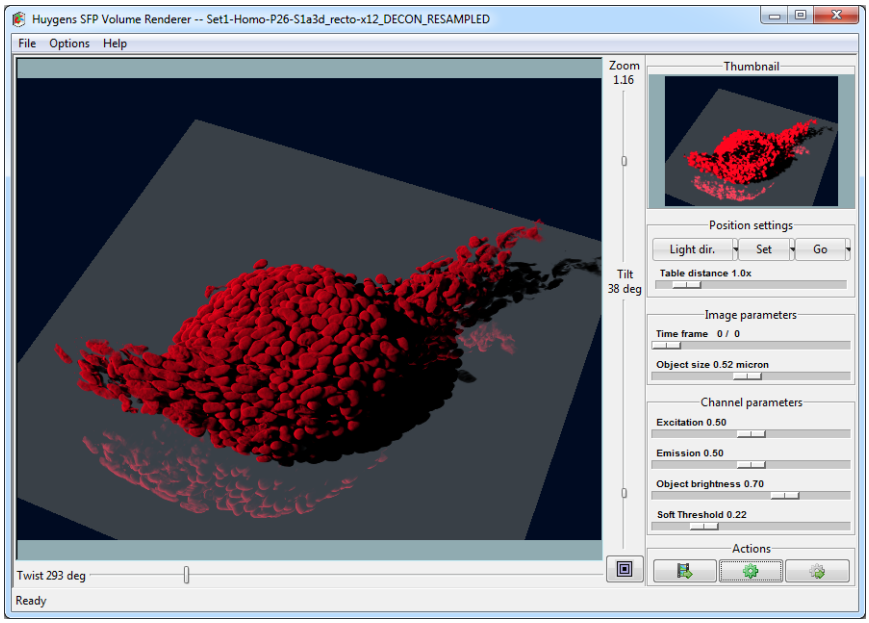

FIGURE 7.2. The SFP Renderer window.

Note that the thumbnail preview (the top right) reflects changes in the configuration instantly, while the large view should be updated manually. To update the large view, press the *fast mode* button ( ) or the *high quality* button ( ).

Threshold

The *Soft threshold* slider in the *Channel parameters* panel at the right affects the threshold level. The application of a threshold is a preprocessing step that *reduces the background* in the image, i.e. voxels with intensity values below the threshold value become transparent. Contrary to a standard threshold, which is 'all or nothing' (values above the threshold are kept, values below it are deleted), the soft threshold function handles images in a different way. It makes a *smooth transition* between the original an the deleted value.

Object Size

The characteristic object size can be set by the *Object size slider* in the *Image parameters* panel at the right. This parameter affects both the excitation and the emission *transparency*. While traveling through the object, the light intensity is attenuated to some degree. This enables us to define some definition for penetration depth at which the light intensity is decreased to some extent, for instance 10 % of its initial value. This penetration depth should be in line with the object size. A transparent object is small with respect to the penetration depth. Thus for the same physical properties of the light one object can be transparent while the other is oblique due to its size. To find a reasonable range in transparencies the object size may be altered. The initial object size is computed from the microscopic sampling sizes and number of pixels the image is composed of. If the microscopic sampling sizes of the image are incorrect, then the object size is set according to some default parameters. and may not be related to the actual object size.

Saving Scenes

Choose FILE→SAVE SCENE... to save the rendered scene as a Tiff file.

Advanced Usage

SFP Fundamentals

The voxel values in the image are taken as the *density of a fluorescent material*. In case of a multi channel image, each channel is handled as a different fluorescent dye. Each dye has its specific excitation and emission wavelength with corresponding distinct absorption properties. The absorption properties can be controlled by the user (See the *transparencies* in Table 7.1 on page 35). The different emission wavelengths give each dye its specific color.

To excite the fluorescent matter light must traverse other matter. The resulting attenuation of the excitation light will cause objects, which are hidden from the light source by other objects, to be weakly illuminated, if at all. The attenuation of the excitation light will be visible as shadows on other objects. To optimally use the depth perception cues generated by these shadows, a flat *table* below the data volume is placed on which the cast shadows become clearly visible. In Figure 7.2 the table is rendered as a mirror.

After excitation the fluorescent matter will emit light at a longer wavelength. Since this emitted light has changed wavelength it is not capable to re-excite the same fluorescent matter: multiple scattering does not occur. Thus only the light emitted in the direction of the viewer, either directly or by way of the semi reflecting table is of importance. By simulating the propagation of the emitted light through the matter, the algorithm computes the final intensities of all wavelengths (the spectrum) of the light reaching the viewpoint.

The properties of the interaction between object and light (transparency), both for excitation and emission, can be adapted interactively by the user to produce different sceneries.

Render Parameters

Table 7.1 gives an overview of all render parameters in the SFP Renderer.

TABLE 7.1. SFP render parameters

Parameter	Description
Light direction	Change the angle for the excitation illumination.
Table distance	Adjust the distance between the object and the table.
Time frame	Set the time frame (in case of a time series).
Object size	Adjust the total transparency of the rendered object. See "Object Size" on page 34
Excitation	Adjust the excitation transparency for the matter in the selected channel.
Emission	Adjust the emission transparency for the matter in the selected channel.
Object brightness	Set the intensity level for the excitation light source for the selected channel.
Soft threshold	Adjust the threshold level for the selected channel. See "Threshold" on page 34.

Render Options

Table 7.2 gives an overview of the different render options that are available through the

TABLE 7.2. Render options for the SFP Renderer.

Option	Description
COMPUTE SHADOW	Enable or disable the computation of shadows.
SHOW SVI LOGO	Show or hide the SVI logo at the bottom right.
SMALL THUMBNAIL	Reduce the size of the preview thumbnail. This enhances the interactivity on slower computers.
TABLE STYLE	The table style can be set to OFF (hidden), NORMAL (grey), and MIRROR. (reflective).
ANIMATION FRAME COUNT	Set the number of frames that will be rendered in a movie. 180 frames with a frame rate of 24 fps result in a movie with a duration of 7.5 seconds.
ANIMATION FRAME RATE	Adjust the frame rate; a rate of 24 frames per second is fine for smooth movies.
VIRTUAL RENDER SIZE	Adjust the size of the rendered image. When the render size exceeds the display area, then use the *middle mouse button* to pick up and move the rendered image.
RENDER QUALITY	Set the default quality (FAST, NORMAL, or BEST). This setting will be used for rendering animations.
COLOR MODE	Choose between EMISSION COLORS or GLOBAL PALETTE (See See "Adjusting the Global Color Scheme" on page 104).

OPTIONS menu. The ANIMATION FRAME COUNT, ANIMATION FRAME RATE and RENDER QUALITY apply to the rendering of simple movies as explained in the next section.

Templates

All scene settings, i.e. both the render options and all parameters, can be exported to a template file via FILE→SAVE SCENE TEMPLATE. The template files have the extension `.hgsv` and they can be applied to any image that is loaded in the SFP Renderer, but keep in mind that the sampling sizes of the data affect the transparency.

Simple Animations

The Huygens *Movie Maker* (See "The Movie Maker" on page 43) allows to create easily sophisticated animations using the MIP, SFP, and Surface Renderer.

Without the Movie Maker the SFP Renderer has the option to make simple animations of the image, changing the view point in different frames. Set the render parameters for the first frame and click SET→HOME in the *Position settings* panel at the right. Now adjust the viewpoint for the final frame, and click SET→END. Also the *frame count*, *frame rate*, or other render options in the OPTIONS menu may be adjusted. Finally press the *animate* button ( ), and select a directory to save the AVI movie or the TIFF frames to.

The exported AVI files use the MJPEG[2] codec and can be loaded in most movie players, including Windows Movie Player and Apple Quicktime. TIFF frames are useful to combine multiple animations or edit the movie in e.g. Windows Movie Maker.

2. `http://en.wikipedia.org/wiki/Mjpeg`

CHAPTER 8

The Surface Renderer

An iso-surface is a 3D surface representation of points with *equal intensities* in a 3D stack; it is the 3D equivalent of a contour line (See Figure 8.3). The Huygens Surface Renderer is a powerful 3D visualization tool that enables the visualization of these surfaces and thus the representation of the data in a convenient way to clearly see *separated volumes*. Shading enhances the perception of 3D shapes and texture (See Figure 8.1).

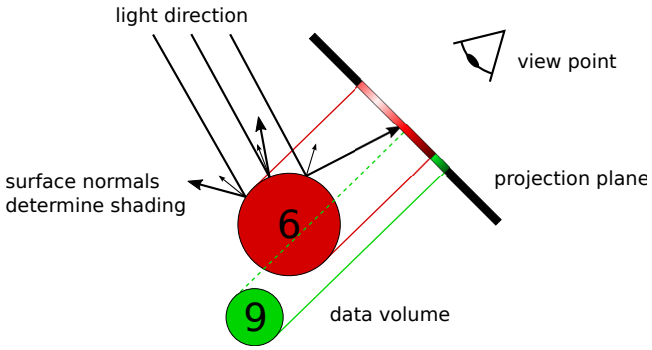

FIGURE 8.1. A schematic overview of surface rendering. The orientation of the surface determines the shading.

Besides iso-intensity surfaces, this renderer is able to generate MIP projections which are blended with the surfaces to be used as a reference to the original microscopic data (See Chapter 6 on page 29).

Because the Surface Renderer is based on *fast ray tracing* that runs efficiently on multi-core computers, there is no need for any special graphics card as would be necessary for conventional polygon based techniques.

To start the Surface Renderer, right-click on a thumbnail and select VIEW→SURFACE RENDERER from the pop-up menu, or choose VISUALIZATION→SURFACE RENDERER from the main menu.

Orientation and Zoom

Adjust the viewpoint by moving the *Tilt* and *Twist* sliders (See Figure 8.2) or by dragging

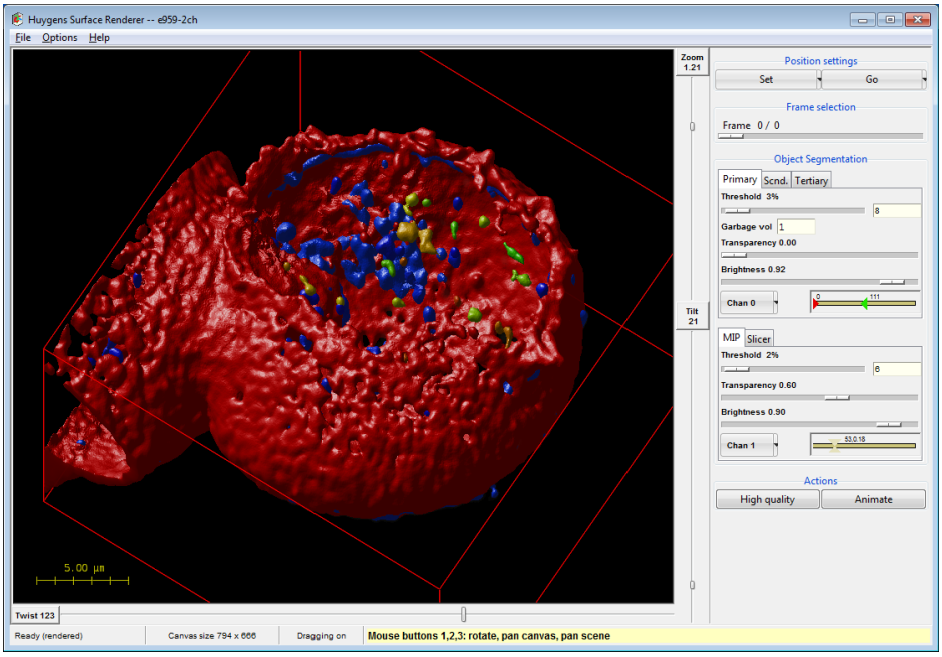

FIGURE 8.2. The Surface Renderer window.

the mouse pointer on the large view. The magnification can be adjusted using the *Zoom* slider or by using the scroll wheel. Use right mouse button to *pan* the center of the projection

Threshold

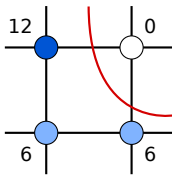

FIGURE 8.3. A contour line for an interpolated value of 5. Because 5 is much closer to 6 than to 12, the distance of the contour to the voxel with value 12 is larger than the distance to the bottom-right voxel with value 6

Use the *Threshold* slider in the *Render pipes* panel to apply different thresholds to the data channels (See Figure 8.3). Voxels that are spatially connected and have intensities above this threshold define *closed volumes*. These volumes are represented by the 3D (iso-intensity) surfaces containing them, each object having a different surface color.

The three *render pipes*, in the *Object Segmentation* frame, referred to as *primary*, *secondary*, and *tertiary*, allow us to define three threshold levels that can be applied to the same or to different data channels. The data channel can be selected using the menu button in the *Object Segmentation* panel. The color range in which the different objects inside a render pipe will be displayed can be adjusted with the hue selector next to it.

Saving Scenes

Press the HIGH QUALITY button in *Actions* panel to apply full scene *anti aliasing* to the rendering and choose FILE→SAVE SCENE... to save the rendered scene as a Tiff file.

Advanced Usage

Adding a Maximum Intensity Projection

Besides the surface pipes there are additional rendering pipes to redirect data to the scene. The *MIP* pipe works projecting the voxels with maximum intensity laying in the path of the rays traced along the viewing direction (See Chapter 6 on page 29). In combination with the surface pipes, very clear representations can be obtained of the different objects in the image. The MIP of a channel can be a good spatial reference for the objects in other channels.

Adding a Slice

The *Slicer* pipe is available to represent a single slice of the 3D dataset in its corresponding location.

Render Parameters

Table 8.1 gives an overview of all render parameters in the Surface Renderer.

TABLE 8.1. Surface render parameters

Parameter	Description
Frame	Set the time frame (in case of a time series).
Threshold (surface)	Adjust the threshold level for the selected pipe, i.e. the intensity for which the iso-intensity surfaces are defined. See "Threshold" on page 38
Threshold (MIP)	Set the soft threshold level for the MIP pipe. See "Threshold" on page 29
Garbage volume	Volumes that contain less voxels than defined by the *garbage volume* parameter will not be rendered. This is useful for rendering only significant objects in noisy images.
Transparency	Set the level of transparency to other pipes.
Brightness	Adjust the brightness for the selected pipe.
Slice Z-position	Set the position of the slice in the *Slicer* pipe.

Render Options

Table 8.2 gives an overview of the different render options that are available through the

TABLE 8.2. Render options for the Surface Renderer.

Option	Description
ANIMATION FRAME COUNT	Set the number of frames that will be rendered in a movie. 180 frames with a frame rate of 24 fps result in a movie with a duration of 7.5 seconds.
ANIMATION FRAME RATE	Adjust the frame rate; a rate of 24 frames per second is fine for smooth movies.
VIRTUAL RENDER SIZE	Adjust the size of the rendered image. When the render size exceeds the display area, then use the *middle mouse button* to pick up and move the rendered image.
TRANSPARENCY DEPTH	This option defines how different surfaces are seen through the others: *Simple*: see through one surface, the surface closest to the viewer. Quite often this is sufficient. *Normal*: see through two surfaces. *Deep*: consider many more screening levels, making the final rendering computationally more complex.
BOUNDING BOX	Enable or disable the bounding box, or adjust the line color.
SCALE BAR	Enable or disable the scale bar.
SHOW SVI LOGO	Hide or show the SVI logo at the bottom right.
High quality MIP	Render the MIP pipe in high quality mode.
CENTER SCENE	Undo both the panning of the projection center (right mouse button) and the rendered image itself (middle mouse button).

OPTIONS menu. The ANIMATION FRAME COUNT, ANIMATION FRAME RATE and RENDER QUALITY apply to the rendering of simple movies as explained in the next section.

Templates

All scene settings, i.e. both the render options and all parameters, can be exported to a template file via FILE→SAVE SCENE TEMPLATE. The template files have the extension `.hgsv` and they can be applied to any image that is loaded in the Surface Renderer.

Simple Animations

The Huygens *Movie Maker* (See "The Movie Maker" on page 43) allows to create easily sophisticated animations using the MIP, SFP, and Surface Renderer.

Without the Movie Maker the Surface Renderer has the option to make simple animations of the image, changing the view point in different frames. Set the render parameters for the first frame and click SET→HOME in the *Position settings* panel at the right. Now adjust the viewpoint for the final frame, and click SET→END. Also the *frame count*, *frame rate*, or other render options in the OPTIONS menu may be adjusted. Finally press the *Animate* button in the *Actions* pane, and select a directory to save the AVI movie or the TIFF frames to.

The exported AVI files use the MJPEG[1] codec and can be loaded in most movie players, including Windows Movie Player and Apple Quicktime. TIFF frames are useful to combine multiple animations or edit the movie in e.g. Windows Movie Maker.

1. http://en.wikipedia.org/wiki/Mjpeg

CHAPTER 9

The Movie Maker

The Movie Maker is a tool that allows to create easily sophisticated animations of multi-channel 3D images using the powerful Huygens visualization tools. Animations from the *MIP Renderer* (See Chapter 6 on page 29), the *SFP Renderer* (See Chapter 7 on page 33), and the *Surface Renderer* (See Chapter 8 on page 37) can be combined in a single movie.

The Movie Maker assists the user in creating the *key frames* that define the main scenes, and the smooth *transitions* between them. Interactive manipulation of the scenes is possible using the interfaces of the renderers or by dragging nodes in the *Timeline*.

The movies can be exported to AVI files or to TIFF series that can be combined with other software. Movie projects can be saved for later editing or for usage with other 3D datasets.

An introductory tutorial can be found in the HELP menu at the top right. This interactive tutorial guides the user step by step through the process of creating a simple movie.

An Overview

Figure 9.1 shows the Movie Maker's user interface. The numbered areas are:

1. The *storyboard*: this filmstrip shows the main elements of the movie, which are the keyframes and the transitions between them.
2. The *preview area*: this mini movie player quickly creates a low-resolution version of the movie.
3. The *timeline*: this interactive plot shows how render parameters change over time. Use the menu-button below this timeline to the render parameter to be visualized. Green nodes, representing render parameters at each keyframe, can be dragged vertically.

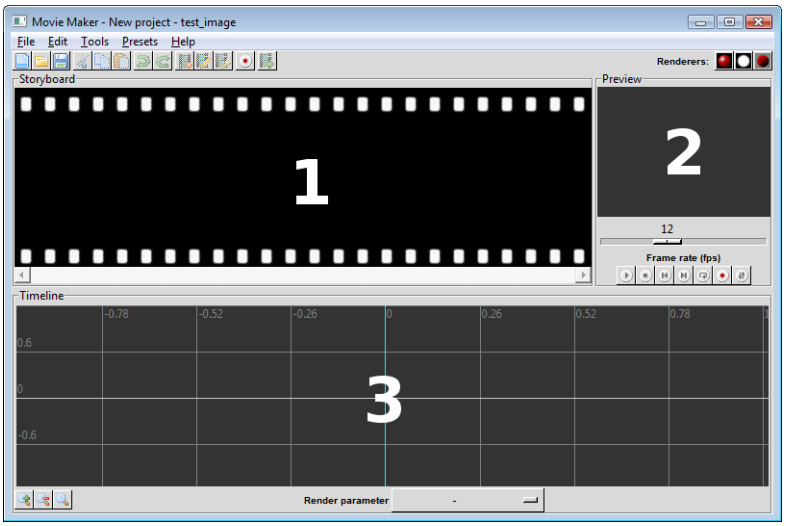

FIGURE 9.1. An overview of the Movie Maker user interface, showing the *storyboard* (1), the *preview area* (2), and the *timeline* (3).

Creating and Adjusting Keyframes

About Keyframes

A keyframe defines a *control point* within a transition. This can be either a start point, end point, or an intermediate point in time. The appearance of the 3D rendered image is *fixed* in these frames. In between keyframes, the Huygens Movie Maker calculates a smooth or linear transition (a technique called *tweening*).

Inserting New Keyframes

To add the first keyframe to the storyboard, one of the renderers should be launched from the Movie Maker window by pressing the corresponding button (, , or). In the renderer a scene can be defined using the available controls; see Chapter 6, Chapter 7, and Chapter 8 for more information on these windows. Once the looks of the first frame are satisfactory, press the *add keyframe* button ( ) to capture this configuration and add the keyframe to the storyboard. All render parameters are captured and can most of them be smoothly animated.

Because a movie needs at least two keyframes (the start and end point of a *transition*), the same renderer should be used to define a second configuration. Once the looks of the second frame are satisfactory, press the *add keyframe* button ( ) again. The storyboard will now show two keyframes with an arrow in between. This arrow indicates transition from one keyframe to another.

The Huygens Movie Maker allows the user to mix keyframes from different renderers, but transitions can only be made between keyframes from the same renderer, as shown in Figure 9.2.

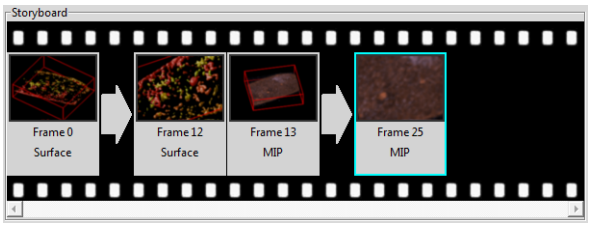

FIGURE 9.2. The Movie Maker's storyboard showing two surface renderer keyframes, followed by two MIP renderer keyframes.

Editing Keyframes

To edit an existing keyframe, double click it or select the frame and choose EDIT→EDIT KEYFRAME from the menu. This will load the keyframe's settings in the corresponding renderer. The renderer's controls can now be used to adjust the 3D scene. To submit the changes to the Movie Maker, press the *add keyframe* button ( ) again. Because the original keyframe is still selected, the Movie Maker will ask if the original frame should be replaced.

Using the Storyboard

Rearranging Keyframes

The storyboard allows the user to copy, delete, and rearrange keyframes. The *cut* ( ), *copy* ( ), *delete* ( ), and *edit* ( ) buttons are activated whenever a keyframe is selected. When there is no keyframe selected, the *paste* button ( ) will append the copied or cut keyframe to the last frame in the storyboard. If one of the keyframes is selected, the Movie Maker will ask where the frame should be inserted. All these operations can be undone by clicking the *undo* button ( ).

Changing Transitions

Transitions can be changed by double clicking the arrow that joins two keyframes. This will pop-up a window in which the duration of transition (in frames or seconds) can be changed, as well as the transition type of the parameters that are animated (See Figure 9.3). For most parameters, the Huygens Movie Maker can calculate linear or smooth transitions. An exception to this is the twist, which also requires a direction of rotation (clockwise or counter clockwise) and a number of rotations.

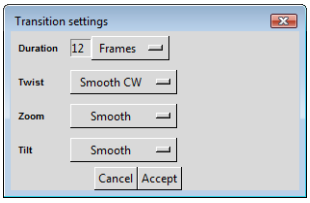

FIGURE 9.3. The transition settings dialog is shown by double clicking a transition arrow in the storyboard.

Playing a Preview Movie

To preview the movie, just press the *play* button ( ) in the *Preview* area (See Figure 9.4). The Movie Maker quickly creates a low-resolution movie and displays it in the preview area. If *loop mode* ( ) is on, the movie will be repeated until the *stop* button ( ) is pressed.

To change the animation settings (*aspect ratio*, *size*, *frame rate*, etc.), press the *render settings* button ( ) in the *Preview* area. The *High quality* setting and the *AVI quality* are not reflected in the movie preview, but only noticeable in the final result.

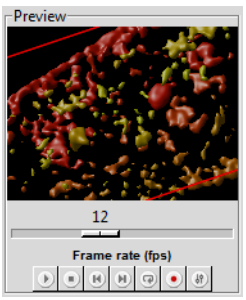

FIGURE 9.4. The preview area.

Export to AVI or Tiff Series

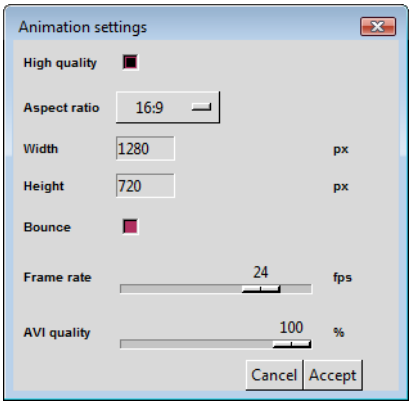

FIGURE 9.5. The animation settings dialog.

If the result is satisfactory, press the record button ( ) below the preview (See Figure 9.4) to render the final movie and export it to AVI or a Tiff series. Before the save dialog appears, the Huygens Movie Maker will show the animation settings dialog (See Figure 9.5), where e.g. the AVI quality can be set. Note that large movies will take several minutes to render.

Once the movie has been exported to AVI, it can be opened in a movie player like Apple QuickTime or Windows Movie Player. To quickly open the last saved movie in the operating system's default movie player, press the button labeled "*Open AVI in external movie player*" ( ).

Working with Movie Projects

Saving Projects

The collection of keyframes, transitions, and animation settings is called a *project*. The image itself does not belong to the project. To save the current project to disk, press the *save* button ( ) or choose FILE→SAVE PROJECT in the menu. The Movie Maker project files have the extension `.hgsa` (Huygens animation template).

The Movie Maker will ask if the project should be saved when one attempts to close the Movie Maker while there are unsaved changes to the project.

Reloading and Appending Projects

If a saved project is reloaded in a Movie Maker that has the same image attached, then the final movie will be exactly the same. However, the Movie Maker allows the user to apply saved projects to different images or append saved projects to the current storyboard. To load or append a project from disk, select FILE→OPEN PROJECT... or FILE→APPEND PROJECT... from the menu respectively.

Besides projects from disk, the Movie Maker has some *presets* that can be appended to the movie. These can be found in the PRESETS menu.

Using the Timeline

Visual Feedback

The timeline is an interactive plot which shows the frame number on the horizontal axis and the value of a render parameter on the vertical axis (See Figure 9.6). It gives a more detailed visual feedback on values of each of the animated parameters. A mouse click somewhere in the timeline area will select the corresponding frame and display a preview of the frame in the preview area. The left and right arrow keys can be used to navigate through the frames.

To zoom in on the timeline, click near the frame of interest and use the scrollwheel or the *magnifying glass* buttons below the timeline to change the zoom level.

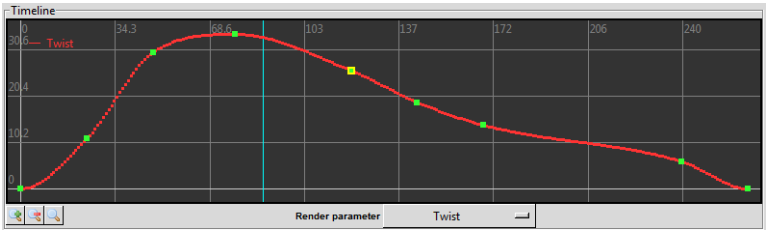

FIGURE 9.6. The timeline area gives a detailed visual feedback on the animated parameters.

Changing Render Parameters

The keyframe nodes are displayed in green and can be dragged vertically to adapt the value of the parameter. The (interpolated) transition frames are displayed as smaller red dots (See Figure 9.6). If a render parameter has been changed, the Movie Maker will recalculate the transitions, update the thumbnails in the storyboard, and update the still of this frame in the preview area.

To select the render parameter which value should be shown in the graph, open the drop-down menu below the timeline and choose the preferred parameter. Because the amount of parameters that can be animated is huge, the menu only lists the ones that change during the movie. If "*Other render parameter...*" is selected, a dialog window will pop up that shows a list of all available parameters.

In order to change one of the parameters in a frame which is not a keyframe, the frame can be converted into a keyframe. To do this, select the frame in the timeline and choose EDIT→CONVERT TO KEYFRAME from the menu.

Advanced Topics

Stretching Movie Length

The number of frames in a transition can be changed by double clicking the transition arrow. However, if a movie gets complex it is easier to use the *stretch tool* to e.g. double the number of frames in all transitions. This tool can be found in the menu TOOLS→STRETCH MOVIE.

The stretch tool shows the *frame rate*, the *number of frames*, and the *duration* in seconds (See Figure 9.7). The two sliders can be used to change the frame rate and the number of frames; changing these will affect the duration of the movie. The stretch tool tries to redistribute the total number of frames over the whole movie in such a way that the relative length of each transition does not change.

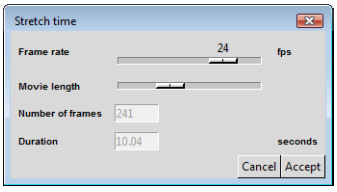

FIGURE 9.7. The stretch dialog helps increasing or decreasing the number of frames in a complex movie.

Synchronizing Transitions in Time and Slice Plane Transitions

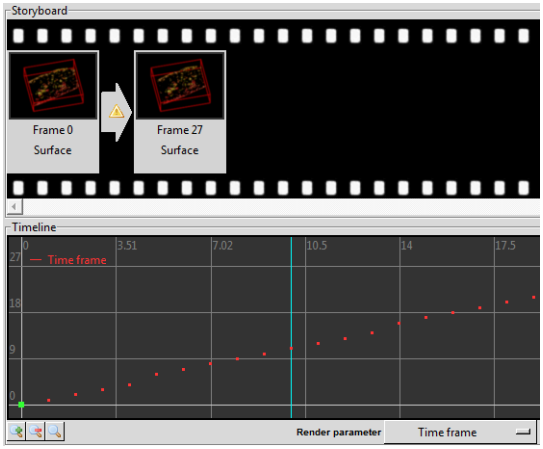

FIGURE 9.8. A transition in which the *time frame* parameter is out of sync with the number of frames. In this example, the transition counts 27 frames, while the *time frame* parameter increases linearly from 0 to 30. Note the irregular change of this parameter in the plot.

Most render parameters, like the *tilt*, *twist*, and *zoom*, can be set to non-integer values. The *time frame* (in case of a time series), and the *slicer z-position*, however, are fixed to integer values. When the transition of such a parameter is not linear, or when the change in value does not match the number of frames, then this parameter is *out of sync*. In that case the Movie Maker will show a warning symbol (⚠) on the transition arrow (See Figure 9.8). In the final result, these asynchronous transitions may show irregularities.

To restore the synchronization, right click on the transition arrow and choose SYNCHRONIZE TIME FRAMES... from the pop-up menu. The Movie Maker will set the transition type to linear and add or remove some frames from the transition to achieve a 1:1, 2:1, 1:2, etc. synchronization.

Creating Loopable and Bouncing Movies

In two simple steps, a movie can be made *loopable*, i.e. it can be played seamlessly in repeat mode:

1. Copy the first keyframe and paste it to the end.
2. Right-click on the final keyframe that has just been pasted, and click SKIP THIS FRAME in the pop-up menu.

The result is the best when all transitions are set to *smooth*.

To create a *bouncing* animation, i.e. an animation that is played in reverse when the last frame is reached, mark the *bounce* option in the animation settings dialog. Doing so will not insert extra keyframes on the storyboard, but the frames are appended in reverse to the final AVI file or Tiff series.

About Movie Quality

In the Huygens Movie Maker, two types of quality can be set:

1. The *High quality* check box determines if the quality of the renderers should be set to the highest possible setting when rendering the final movie.
2. The *AVI quality* scale bar determines the compression level of the AVI file. Set to 100 %, the quality is the best, but then the file size will be large.

These quality settings are not reflected in the movie preview, but only noticeable in the final result.

CHAPTER 10

Introduction to the Object Analyzer

The interactive Object Analyzer (OA) tool allows you to obtain *statistics of individual objects* by clicking on them, or analyzing all objects with a single button press.

In this context, an object is a distinct group of interesting voxels that are spatially connected one to another. Interesting voxels are distinguished from the background by using a seed and threshold criterion. Therefore, defining objects in an image implies:

1. *Segmentation*: Separating interesting voxels from the background according to a given criterion;
2. *Labeling*: Grouping them together and assigning them a distinct name or label.

This is done interactively by the Object Analyzer. To remove too small objects in an early stage from the analysis, a garbage level can be set below which objects are discarded. After that, detected objects are automatically labeled and sent to a continuous iso-surface renderer (See Chapter 8 on page 37).

The Object Analyzer is an extended optional tool, and is enabled by a special flag in the Huygens license string.

This chapter is written in the form of a step-by-step, introductory tutorial to the basic functions of the Object Analyzer. A reference guide that describes all the components of this tool can be found on page 65.

Starting the Object Analyzer

- Launch Huygens Essential or Huygens Professional.
- Load an image you want to analyze. To explore all the OA possibilities, better use a multi channel image.
- Select the image thumbnail, and in the top menu go to ANALYSIS→OBJECT ANALYZER ADVANCED.

You can find this introductory tutorial on-line in the SVI support Wiki[1], from where you can also download the test image we will use in the following explanations.

1. `http://www.svi.nl/FrontPageObjectAnalyzerTutorial`

Please explore the image with the Twin Slicer (See Chapter 4 on page 19) to get an impression of it. This is a deconvolved image, to reduce noise and blur artifacts. It is always a good idea to perform object analysis with deconvolved datasets.

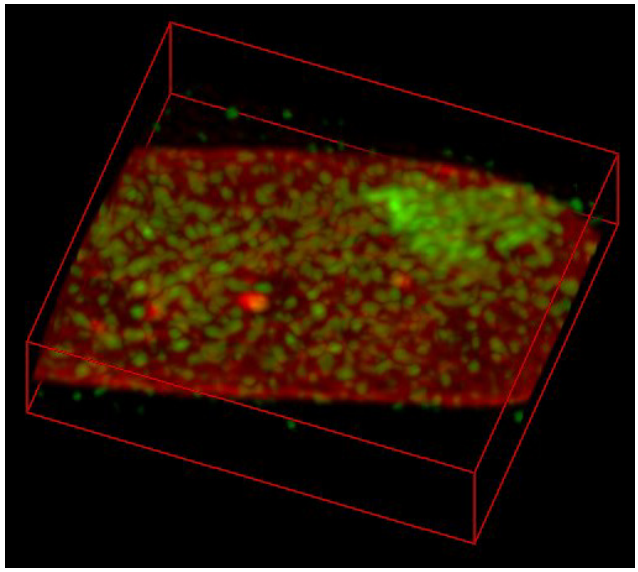

FIGURE 10.1. Maximum intensity projection of the test image. Cell nucleus FISH-stained, recorded at the Nuclear organization Group, SILS, University of Amsterdam (head: Prof. Roel van Driel), under the 3D-Genome research project.

The image in Figure 10.1 is a MIP projection of the test 3D dataset. Notice that there are a few bright objects in the red channel against a more or less homogeneous background (it is actually a quite flat cell nucleus), and lots of scattered objects of different sizes and intensities in the green channel. The red channel is in general very dim with the exception of the inner bright objects and some increase in intensity in its periphery, making something like a border.

When the image is opened it in the OA this is what we first see what is shown in Figure 10.2.

The bright objects in the red channel are recognizable, now from a top view. The image was automatically rendered with some default settings. The intensity range was explored to set a convenient threshold to segment the objects in the first (red) channel, that are shown as iso-surfaces. Every independent object gets a different color, ranging from red to green.

How were these objects separated from the background, isolated from the rest of the intensities in the image so that they could be represented on the rendering as independent entities? They are said to be *segmented*.

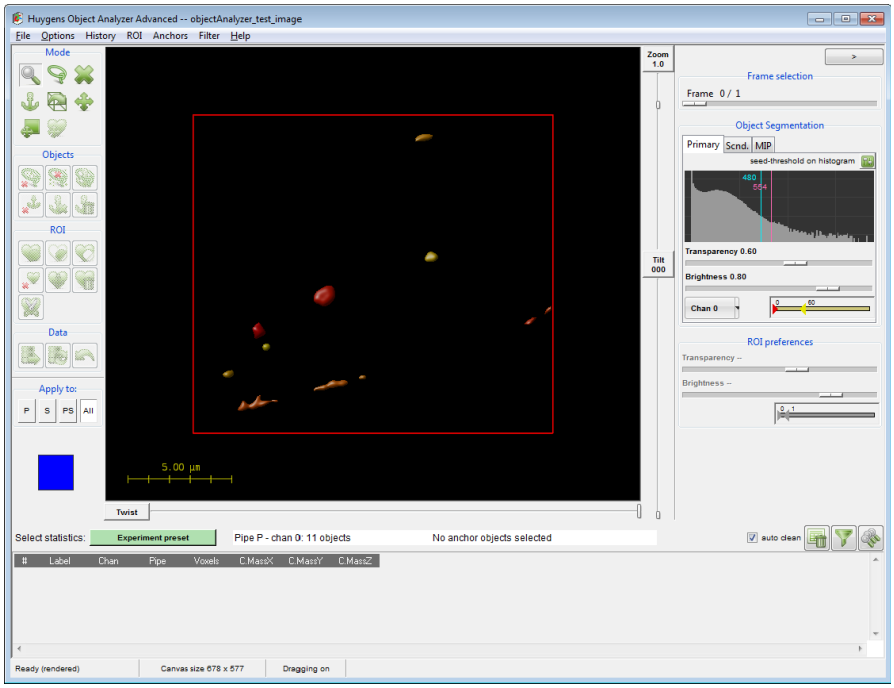

FIGURE 10.2. Start up screen of the Object Analyzer. The objects in the image's first channel are automatically segmented with some default parameters calculated based on the intensity distribution.

Segmenting the Objects: Setting the Threshold

The segmentation method currently available is the Object Analyzer is more flexible than simply setting a threshold: it uses a combined seed-and-threshold method. You can ignore the seed if you want, for starters. It acts as a secondary threshold level, so that objects that do not reach it in intensity (in at least one voxel) are discarded. In this introductory tutorial we will basically ignore the seed, an leave it at the automatic linked value, slightly larger than the threshold itself. For more details on how it works, you may refer to the expert on-line tutorial[2].

Let's try different segmentation parameters and see what happens. First we lower the threshold value from the automatically calculated value down to something around 360 (the mentioned numerical values refer to this particular test image, of course). You can drag the blue line in the channel histogram and shift it to lower values, or click on the blue-font label showing the threshold value to enter any number. The threshold line on the histogram can be found at the right of the window, in a pane labeled *Object Segmentation*. Next to it, in magenta, there is the seed level. As it is linked to the threshold by default, when you shift the blue line the other one goes with it.

(You can also switch to a percentage representation of threshold and seed by clicking the small button at the top-right of the histogram. In this alternative slider view, the only one before Huygens 3.5, you can also control whether the seed is linked to the threshold or it remains independent).

2. `http://www.svi.nl/FrontPageObjectAnalyzerExpertTutorial`

We see many things happening when we shift the threshold to lower values. First, the objects we had defined grow in size because more voxels around them get attached to them. Some objects that were separated before grow so much that the now get connected, and they define new single objects. We also get more and more objects of all sizes. To appreciate the details better, you can increase the zoom factor up to 1.30 to see what is shown in Figure 10.3.

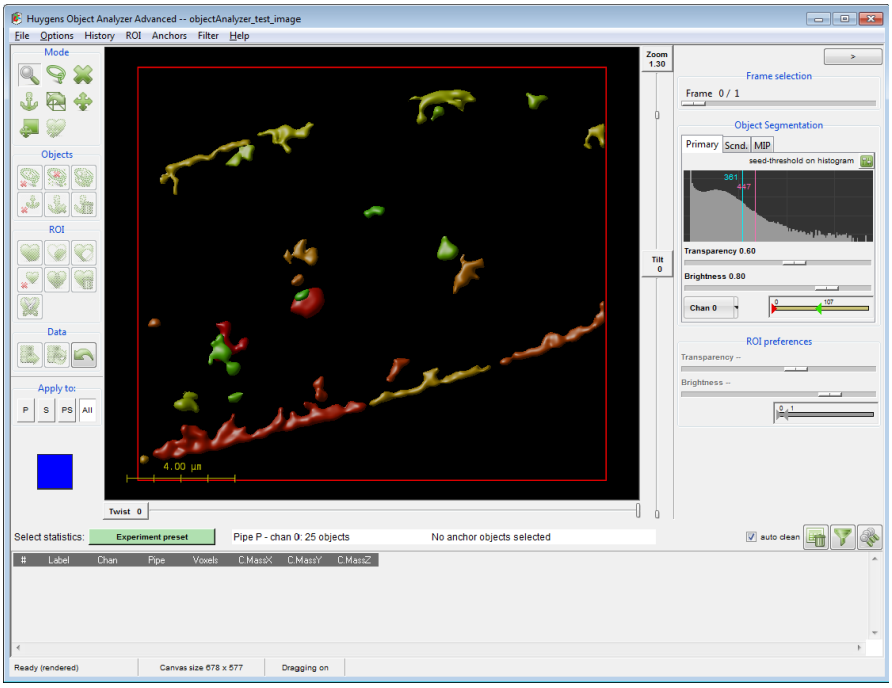

FIGURE 10.3. Different segmentation conditions. Lowering the threshold increases the objects size, and adds more objects to the scene.

In this new segmentation, after lowering the threshold, we have made the original objects larger, but also included many new ones that may be in the way. If you are not interested in all those new objects now, you can try to remove them. Here is where the seed plays an interesting role, but we decided to forget it by now. We can also filter the objects based on their measurable properties, or just use a garbage voxel level, to discard objects that are too small.

Setting a Garbage Volume Level

A quick way of removing disturbing objects is the GARBAGE VOLUME. You can find this entry in the alternative slider view of the threshold, to which you can switch by clicking the small button ( ) at the top-right of the histogram. The garbage entry will be shown right below the seed slider. Objects with a number of voxels below the garbage level are discarded. This means that when you set it to 1, no segmented object is discarded, but if you set it to e.g. 100, any object with a volume smaller than 100 voxels will be removed.

So far for the segmentation. You can apply some post-segmentation filtering in complex ways, the details are explained in the expert tutorial. In the top menu you can also find some predefined Filters that you can use for quick access to discard objects based on some basic geometrical properties or on the way they relate to other objects.

By now, these two tools (threshold and garbage level) are already powerful enough to continue our exploration of other Object Analyzer features.

Interaction with the Objects

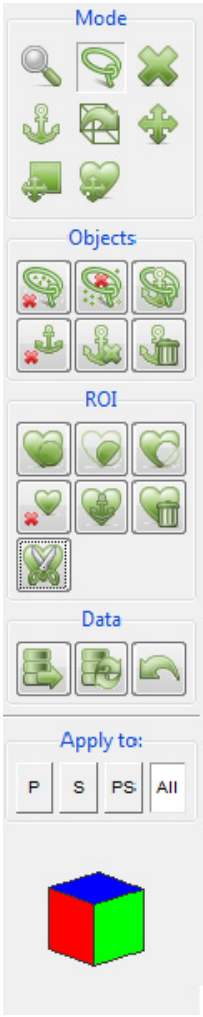

On the left of the window you can find a column of buttons. Many of the buttons are deactivated at this moment, they will be activated when it makes sense later. But almost all in the first group are always available. They control the *Mouse mode*, i.e. what the mouse does when you left-click on the image. When you hover with the mouse over these buttons you get a tip on what they are for. If they have a keyboard shortcut to activate them, it is also shown here. For example, the first mouse mode is intended to Analyze objects, and you can always activate it from within the rendering view by pressing the keyboard key '`1`'.

These are the currently available mouse modes (from left to right, and from top to bottom):

- **Analyze object** lets you click on different defined objects and obtain the local statistics.
- **Select area** lets you define a 2D region on the current view of the image so you can do different things with it: analyze or discard objects below it, anchor them as references, or interactively define regions of interest (ROI). We will see what all this is useful for.
- **Discard object** lets you discard irrelevant objects one by one. Just select this mode and click on the disturbing objects.
- **Select object** (as anchor) lets you select and deselect objects to be `anchors', for example to act as references to measure distances from other objects when asking for local statistics. When you set an object as a reference anchor it will `light up' and change color on the screen to indicate its new status. It is possible to select a group of anchor objects, and you can operate with them through the Anchors menu.
- **Rotate scene** interacts with the full image to rotate it in the space, by dragging the mouse pointer on the rendering view. That can also be achieved by moving the Tilt and Twist sliders along the rendering.
- **Pan scene** interacts with the full image to move it in space *laterally*. This means that you can pan the scene in the 2D plane of your screen, not along the third dimension, along your line of sight.
- **Pan canvas** is similar to 'pan scene' but not exactly the same. It allows you to explore the canvas by not re-rendering the scene. This only makes sense when you have a canvas larger than your rendering window, of course. See OPTIONS→VIRTUAL RENDER SIZE in the top menu.
- **Shift the ROI**. This is the only mouse mode that is not always enabled: you need to have defined a *region of interest* (ROI) before moving it around.

The buttons in the second, third and fourth groups are not mouse modes: they do not set new behaviors for your mouse when interacting with the objects view but execute operations on previously defined conditions, for example deleting the objects under a selected

area, or discarding every object that has not been selected as anchor. We will not consider them in this basic tutorial, as you don't need them to perform basic measurements! Please see the "Object Analyzer Component Reference" on page 65 for a detailed description of these toolbox buttons.

At the end of the buttons column there's a colored reference cube that will help you in orienting in space when you rotate the dataset, specially with large zoom factors that do not let you see the surrounding box frame in the rendered image. The initial view of this cube is the blue top face corresponding to $z = 1$. Hovering the mouse over the cube faces brings a tooltip with the face label (x, y and z with values 0 or 1).

The *Analyze object* mouse mode is the default one. Click on an object and see what happens. You probably already did by now and noticed that some parameters were reported on the table at the bottom of the window. We will see how to report even more in a minute.

First, we need to briefly explain what a pipe is, just to know how to select what channels from the image you want to analyze.

Render Pipes

When we opened the analyzer we got the first channel of the image directly shown on the screen, but we can change that and select the other channel from the multi-channel original image to explore it. For that we simply need to select another channel to be shown in the PRIMARY PIPE: where it reads *Chan 0* you just select *Chan 1* (See Figure 10.4).

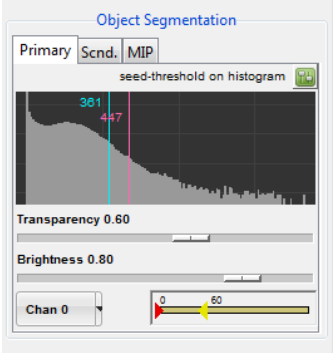

FIGURE 10.4. Rendering pipes control panel.

But what if you want to explore both channels at the same time?

The term *pipe* suggests that data go from your original image to the final rendering in the Object Analyzer through a 'computational duct' in which some processing occurs. You have two of these pipes in the OA to redirect data through.

So we have a secondary pipe too. You can activate its control pane by clicking on the tab that reads SCND. 'Secondary' does not mean here 'less important', it is just that we have the first pipe (*Primary*) and the second pipe (*Secondary*), abbreviated P and S in some places. We could have called them also *Red* and *Green* pipes, independently of their real colors, like it is common in conventional colocalization analysis (and the name may so be changed in future versions of the software). We keep numbers to refer to image chan-

nels, which is something different, because in each of these pipes you can put any channel you want, in any order. You can even segment the same channel twice, with different conditions in each of the pipes.

All the different objects in a pipe are colored differently to be able to distinguish them. The range of colors assigned to each pipe can be controlled with a HUE SELECTOR (See "Hue Selector" on page 105). You can collapse this range completely if you want that all objects in a pipe get the same color.

There is also a maximum intensity projection (MIP) pipe that doesn't interfere in the analysis but that can be used to set a spatial reference to our eyes.

Before continuing our exploration, let's put channel 0 again in the primary pipe with a threshold of 40% (488 in absolute terms), nothing (Off) in the secondary pipe, and channel 0 again in the MIP, so we see something like Figure 10.5.

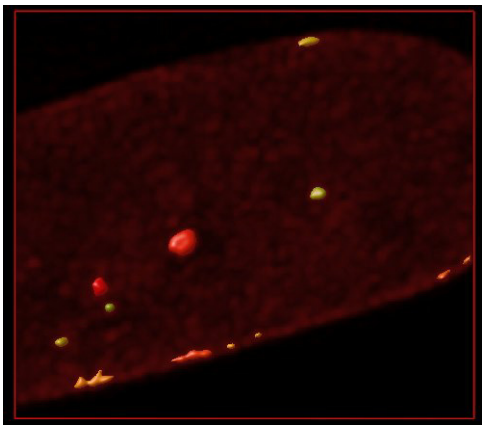

FIGURE 10.5. Objects and MIP. Objects in a surface pipe rendered together with a MIP pipe for spatial reference. The data channel is the same one in both pipes.

Object Statistics

The Object Analyzer can measure a lot of different parameters on the segmented objects, but only a few are reported by default. Otherwise the statistic table would be too saturated with information and you will not be able to find your way in it. To obtain object information you can click on each object interactively while you are in the *Analyze object* mouse mode, or press the ANALYZE ALL button on the top-right of the table to automatically process all the segmented objects in all pipes. (When you have an area selected on the screen, enclosing a few objects, this button analyzes only these objects).

There are many parameters you can measure for each object, not only about its `internal' information but also about its relationship with other objects and regions in the image. Let's see now how to retrieve all this information.

Configuring the Reported Parameters

To simplify the usability of the OA there is a big button next to *Select Statistics* on the top left of the table that reads EXPERIMENT PRESET at start-up. Click on it and you will get a preset selection dialog like the one in Figure 10.6.

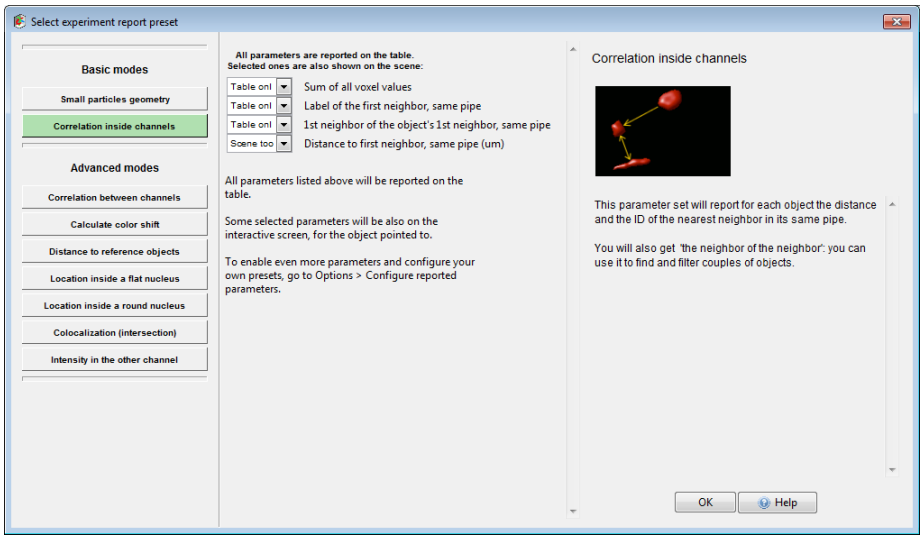
Select experiment report preset
Basic modes
Small particles geometry
Correlation inside channels
Advanced modes
Correlation between channels
Calculate color shift
Distance to reference objects
Location inside a flat nucleus
Location inside a round nucleus
Colocalization (intersection)
Intensity in the other channel
All parameters are reported on the table.
Selected ones are also shown on the scene:
Table onl
Sum of all voxel values
Table onl
Label of the first neighbor, same pipe
Table onl
1st neighbor of the object's 1st neighbor, same pipe
Scene too
Distance to first neighbor, same pipe (um)
All parameters listed above will be reported on the table.
Some selected parameters will be also on the interactive screen, for the object pointed to.
To enable even more parameters and configure your own presets, go to Options > Configure reported parameters.
Correlation inside channels
This parameter set will report for each object the distance and the ID of the nearest neighbor in its same pipe.
You will also get 'the neighbor of the neighbor': you can use it to find and filter couples of objects.
OK
Help

FIGURE 10.6. The Experiment Presets dialog. Collections of parameters can be selected to be reported on the table and interactively on the scene, depending on the experimental needs.

On its left column a series of different experimental needs are listed. When you click on each of them, a new list of parameters is listed in the middle column, and a description is shown on the right column. Even more, when you hover with your mouse over the listed parameters you get a tooltip text explaining each parameter with more detail. Please read the descriptions briefly to get an impression of them. Let's select now the preset called CORRELATION INSIDE CHANNELS and inspect it in detail. Please read the descriptive text.

"This parameter set will report for each object the distance and the ID of the nearest neighbor in its same pipe."

For each object we are interested in, we are going to find also the closest object.

Notice that from the listed parameters in this preset only one is selected by default to be reported on the scene. All of them will be calculated and reported on the table, but only the selected ones will be shown on the rendering canvas as well, when you interact with the objects. In this case it is smart to report only the distance there.

Once you have selected this CORRELATION INSIDE CHANNELS preset, press OK to continue.

Measuring the Objects

Make sure you have the *Analyze object* mouse mode selected, click on a nice object on the screen, and notice the new columns that appear on the statistics table.

Before showing the actual object details, a row in the table informs you about the segmentation conditions for this pipe. Then comes the data itself. When you move your mouse over the column titles at the top of the table you get a description of each parameter at the very bottom of the window. Distance to first neighbor, same pipe (um) As many parameters are available, sometimes measuring things complex to describe, there is only room for a cryptic label in the column title. You can always find out what each parameter is by looking at its tooltip. The description of each parameter will be also stored in your file when you export the table later.

The tooltip of the column header also shows the sum and the average of all the cells below it, that you can optionally restrict to include data of a given pipe only.

Click on another object to add its data to the table. You will see that the conditions are not reported again, because they have not changed. A checkbox option at the top-right of the table clears it automatically when the segmentation or report conditions change, so you always have an organized table. Deselect this option to simply accumulate rows on the table, so you can save it later and process it with another program. Another button next to that checkbox clears the table manually whenever you want.

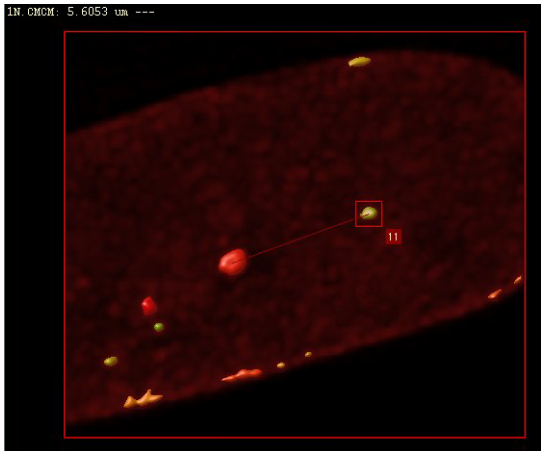

FIGURE 10.7. Interaction with the objects. When the distance to the nearest neighbor is computed, it is also displayed on the screen for the object aimed at.

You may also have noticed that something else happens when you interact with the rendered objects. In the example of Figure 10.7 the user clicked on object 11. A line joins the center of mass (CM) of this object with the center of mass of the nearest neighbor, and this distance is reported at the top of the window. (This is the only parameter reported on the canvas because that is the only one selected on the preset selection dialog, but all the interesting information is now included in the table too).

Other Measuring Parameters

There are two basic parameter presets: this one we have used here to explore the nearest neighbors, and another one to retrieve morphological parameters about objects, called SMALL PARTICLES GEOMETRY.

This other parameter preset can report object information like *length*, *width*, *aspect ratio*, and *sphericity*. Details on how these parameters are defined can be found on page 61.

The presets are organized in basic and advanced modes. The basic modes refer to presets on which you don't have to do much before starting measuring: you click on the objects and you get the desired information. The advanced modes require that you define some conditions first: to define a ROI, for example. They are not intrinsically more complex, but they demand from the user to know how to set these reference conditions first.

For more details about any parameter preset just click the Help button you can find at the lower-right part of the preset selection dialog, and follow the on-screen tooltips during the interaction with the module.

Exploring the Table

Let's try another thing: click on the ANALYZE ALL button at the top-right of the table so the data of all objects are gathered. When the computations are finished (it should be quite fast in this example) move your mouse over the table rows and see what happens. The object corresponding to the current row will be highlighted on the canvas, and the distance to its nearest neighbor will be shown.

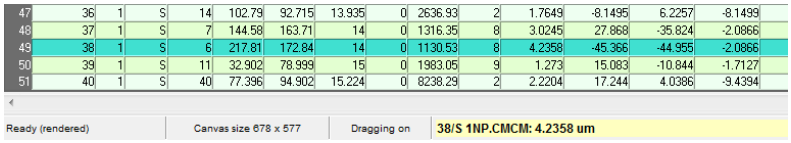

47	36	1	S	14	102.79	92.715	13.935	0	2636.93	2	1.7649	-8.1495	6.2257	-8.1499
48	37	1	S	7	144.58	163.71	14	0	1316.35	8	3.0245	27.868	-35.824	-2.0866
49	38	1	S	6	217.81	172.84	14	0	1130.53	8	4.2358	-45.366	-44.955	-2.0866
50	39	1	S	11	32.902	78.999	15	0	1983.05	9	1.273	15.083	-10.844	-1.7127
51	40	1	S	40	77.396	94.902	15.224	0	8238.29	2	2.2204	17.244	4.0386	-9.4394

FIGURE 10.8. Exploring the table. Hovering the cursor over the table rows also highlights the corresponding object on the scene.

A good way to find an object in a very long table is by clicking on it on the rendering canvas while the *Analyze object* mode is active: the table will be shifted to show its corresponding row, and it will be highlighted.

By moving the mouse over the column titles you get an expanded description of the parameter plus some basic statistics of all the cells in that column. You can also right-click on a column and select STATS for more detailed statistics.

To plot a histogram of the distribution of values in a column, select the column (or a subset), click with the right mouse button, and select HISTOGRAM from the pop-up menu.

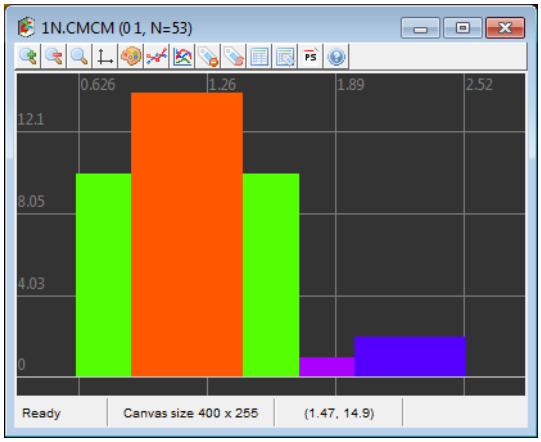

FIGURE 10.9. Histogram representation of one column values. After computing all objects sphericity parameters, for example, and having them reported on the table, we can plot a histogram of the particle sphericity distribution.

Storing your Results

In the top menu you can do FILE→SAVE OBJECT STATISTICS to export the table you see on the screen to a file that you can import elsewhere. The file can be read directly, or imported in conventional spreadsheet programs or data plotters. The table can be stored as a plain text file, a csv-file, that can be imported in e.g. Microsoft Excel, OpenOffice Calc or GNUplot, and as an m-file, that can be imported in Matlab, to do further analysis and/or calculations.

You can also save the current scene as you see it on the rendering canvas to a TIFF file on disk with FILE→SAVE SCENE. In the OPTIONS menu you can find different options that affect how the scene is rendered. You can set the MIP pipe to high quality, for example, or render each analyzed object together with its numeric ID label, so that you can link them with the exported data visually.

In the HISTORY menu, you can save your analysis history as a template, to reproduce it using the same or other dataset.

Further Reading

This tutorial has covered the very basic features to the Huygens Object Analyzer. If you want to learn much more, consider following the expert on-line tutorial[3], where you will be told on many other powerful things you can do with this versatile analyzer.

A reference description of the Object Analyzer components can be found in Chapter 10 on page 49.

3. `http://www.svi.nl/FrontPageObjectAnalyzerExpertTutorial`

CHAPTER 11

Object Analyzer Geometry Measurements

Iso-surface

In the Object Analyzer, the threshold set for the segmentation is also used to define an iso-surface around the object.

An iso-surface is a 3D surface representation of points with equal values in a 3D data distribution. Is the 3D equivalent of a 2D contour line (See Figure 8.3 on page 38).

Based on this, one can measure volume and surface in *high resolution*, by fine polygonization at a sub-pixel level.

Principal Axis

Segmented objects are geometrically analyzed in terms of their principal moments of inertia. (In this sense, the recorded light intensity registered in the image is used as density: the pixel 'values' are interpreted as local mass, so brighter regions weight more).

The principal axes of an object establish a natural system of reference based on its mass distribution. When you rotate an object around one of its principal axes, the angular momentum is parallel to it. This does not happen in general, and is what makes these axes so special.

Around these axes the principal moments of inertia of the object are defined. For one of these axes, the rotation inertia of the object is minimal (around this axis the object would rotate with the least effort). This axis usually lies along the length of the object. The other two axes are orthogonal to it, and orthogonal to each other.

Length and width

One can easily define a box, with dimensions *L*, *pBoxW0*, *pBoxW1* in the system of reference of the principal axes, that encloses the object completely. The sides of this box are in general not parallel to the main planes of the image, because the principal axes do not coincide with the image *x*, *y*, or *z*-axes in general. It is as if the principal box is rotated with respect to the image in order to properly enclose the object, which may not be aligned with any of the image axes.

The length of the object is the largest distance measured along the three principal axes, it coincides with the largest dimension of the principal box *L*.

One could use the other two dimensions of the principal box as width and thickness of the object, but for some practical uses this may be too simplistic.

The width of the objects is actually computed with a search algorithm that acts as a virtual caliper held perpendicular to the length axis. To find the largest width of an object one would rotate the caliper around the object and repeat this procedure while sliding the caliper along the length axis. However, because microscopic data, even when it is deconvolved, often shows orientation dependent imaging due to the lower axial resolution, structures are often elongated in the axial direction. Moreover, in most cases, the voxels themselves are much higher than they are wide, causing all small objects to be elongated.

Clearly, without correction, the anisotropy in resolution would result in an overestimation of the width. To avoid this problem the rotation angles at which the caliper is held are divided in axial directions and lateral directions.

To measure in the axial directions several slices are taken out of the object and analyzed one by one. Each of these slices is parallel to the 'caliper plane', perpendicular to the length axis of the object. In any slice there are directions more oriented towards the optical axis (axial directions) than others. The largest axial width of each slice is obtained by holding the caliper in these directions and searching for the largest among them.

After all slices along the length axis have been examined, the largest axial width of the object is reported as *WiAx*.

The lateral directions in the caliper plane are the directions closer to the *xy*-plane. The caliper measures now the width of each slice in directions near-perpendicular to the optical axis. Taking the largest figure among these lateral directions might again introduce a bias due to elongation, so now two values are computed per slice: the largest and the smallest width along the lateral directions.

After exploring all slices some global relevant figures are reported. The largest lateral width obtained while sliding the caliper plane along the length axis is reported as *WiLat*. The largest of the smallest widths is reported as *WiLatC*. In case of small objects and moderate to high ratios between the axial and lateral resolution this last value is likely to suffer least from the orientation dependent imaging.

The waist (the smallest of the smallest widths) in the lateral directions is reported as *WaistLat*.

What about objects with a vertical length axis? In that case the 'caliper plane', perpendicular to the length axis, will be horizontal, parallel to the *xy*-plane. In a horizontal caliper plane all directions in it are perpendicular to the vertical *z*-axis. As a result there is no 'most axial direction' in that plane. The software then orients the lateral width towards the *x*-axis and the axial direction towards the *y*-axis.

Sphericity

The sphericity is reported in two ways in the Object Analyzer:

1. The axial sphericity is defined in general as the ratio of the volume of an ellipsoid with axes length *L*, width *W* and thickness *T* to the volume of a sphere circumscribed around the segmented object, defined by its length.

   Because it is based on three axes, it gives an idea of the 3D aspect ratio of the objects.

Depending on what of the previously reported parameters we choose for *W* and *T* we have different practical definitions of axial sphericity.

Probably the most intuitive one is the axial sphericity of the principal box *axSphPB*: the ratio of the volume of an ellipsoid with axes *L, pBoxW0, pBoxW1* to the volume of a sphere circumscribed around the principal box (see above) using the largest side (the length of the object *L*) as diameter (Equation 1).

$$AxSphPB = \left(\frac{pBoxW0 \cdot pBoxW1 \cdot L}{L^3}\right)^{1/3} \quad \textbf{(EQ 1)}$$

2. Another definition (reported as *AxSph*) involves the lateral and axial widths discussed above, obtained with the virtual caliper algorithm (Equation 2).

$$AxSph = \left(\frac{L \cdot WiAx \cdot WiLatC}{L^3}\right)^{1/3} \quad \textbf{(EQ 2)}$$

The roughness sphericity characterizes the roughness of the iso-surface, it measures how close the volume-to-surface ratio is to the one of an ideal sphere. This is conventionally defined as:

$$SurfSph = \frac{\pi^{1/3} \cdot (6V_i)^{2/3}}{A_i} \quad \textbf{(EQ 3)}$$

where V_i is the iso-volume and A_i is the iso-surface of the segmented object. Both sphericity values become 1 for an ideal sphere. The *SurfSph* is the inverse of the 'surface factor' *fs* used in Goetze et al.[1]

Aspect Ratio

Again, the aspect ratio of an object can be defined in terms of different dimensions:

- The axial aspect ratio *AxRatio* is the ratio of the object length to its axial width *WiAx*.
- Similarly, the lateral aspect ratio *LatRatio* is the length divided by *WiLatC*.

More Parameters and Filtering

Many more parameters, geometrical and of other kinds, can be calculated, reported, and used to filter the data by the Object Analyzer. Please refer to the on-line tutorials to learn how to use the different parameters and Experiment presets.

Parameters are available that report:

- *Correlation* inside and between channels, by analyzing neighbor objects
- Location of objects with respect to *reference objects* (anchors).
- Location of objects inside *regions of interest* (for example bodies inside a cell nucleus)
- *Colocalization* of objects (by computing the volume and the intensity overlap of segmented objects in different pipes).

Any of the calculated parameters can be used to filter out objects and further segment your image in elaborated ways.

1. Goetze et al., *Molecular Cellular Biology* 27, p. 4475-4487 (2007)

CHAPTER 12

Object Analyzer Component Reference

This section describes the components of the Object Analyzer (OA) interface.

This section is intended as a quick reference. See “Introduction to the Object Analyzer” on page 49 to learn how to use the basic components in context.

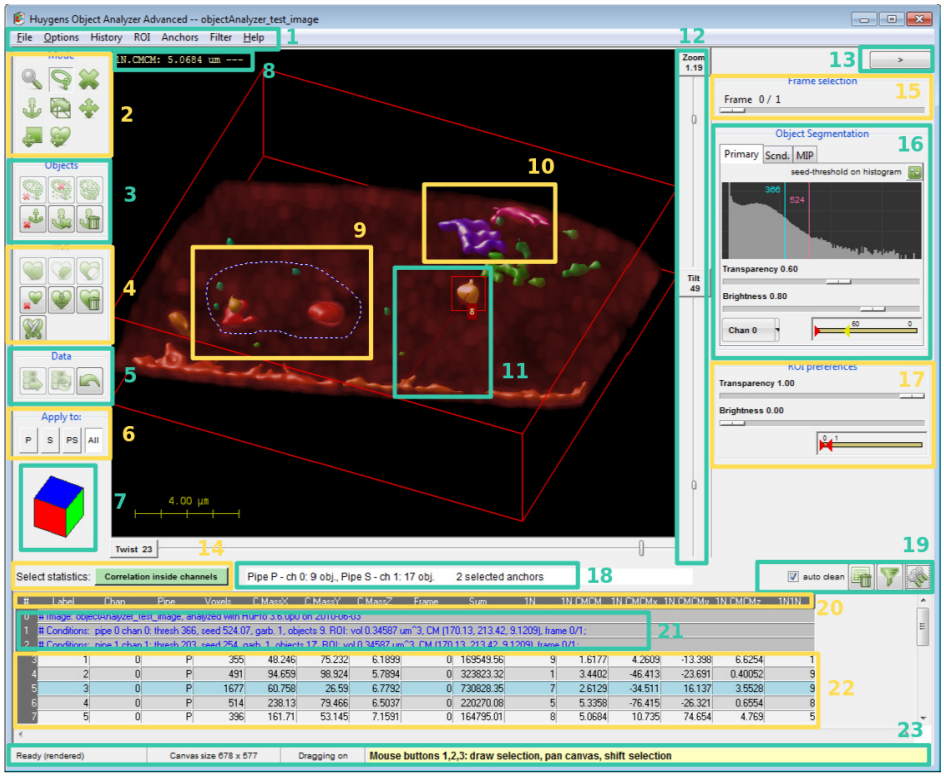

FIGURE 12.1. The Object Analyzer interface. Different regions are enumerated to describe them in sections of this chapter.

Main window components

The OA main window is shown in Figure 12.1 with most of its components in an active state. Different regions of the interface are enumerated to describe them conveniently along this chapter.

When the OA is launched, not all the buttons in the toolbox are enabled. Most buttons are automatically enabled when they are usable, depending on conditions set by the user.

Main Menu (1)

FILE: Entries to save the rendered image, export or clear the object statistics, and analyze all time frames in a series.

OPTIONS: Rendering options and advanced statistics configuration:

- Virtual render size: sets the size of the canvas on which the Ray Tracing algorithm renders the image. This canvas can be larger that the OA window or even the screen (that's why you can pan the canvas to inspect other regions). Like that you can render and save high-resolution TIFF images.
- Transparency depth: controls the number of surfaces considered by the renderer in order to show inner cavities and objects inside objects.
- Bounding box: shows or hides the reference 3D bounding box
- Scale bar: shows a scale bar on the scene. The distances are calculated based on the voxel sizes in the original image microscopic parameters.
- Show ID labels: render the scene showing the numerical ID label of each of the analyzed objects.
- Show SVI logo.
- High quality MIP: enables or disables the high quality rendering mode for the MIP pipe
- Show reference cube (See "Reference Cube (7)" on page 71).
- Show on-screen tooltips for interactive actions.
- Relaxed selection: when active, objects partially outside the selected 2D area are also considered.
- Center scene: moves the point of view to show the center of the dataset.
- Configure statistics report: shows all available parameters to let you select which ones are calculated and reported on the table. This is intended for advanced users, beginners should better use the *Experiment presets* (See "Experiment Presets (14)" on page 73).

HISTORY

- Undo the last operation, or Redo it again.
- Reload original data after cropping the image or discarding objects.
- View the whole operations history up to the current point.
- Load and save analysis templates, so that the current view and parameters can be stored and retrieved, or a whole operations history re-executed with other data. This is also useful to store your analysis steps and reproduce them.
- Set analyzer as in any other open instance of the tool.

ROI: operations to define a region of interest in complex ways, and to modify and save the currently defined ROI.

- Set:
  - Set to extruded selected area: uses the interactively defined 2D area to set the ROI to the volume below it.
  - Set using MIP threshold: use the threshold and data channel of the MIP pipe to set a 3D ROI
  - Make coincide with objects: use the currently segmented objects to define a 3D ROI. Objects from the Primary, Secondary or both pipes can be used depending on the active pipes and the pipe mode (See "Active Pipe Mode (6)" on page 70).
  - Make coincide with anchors: use the currently selected anchors to define a 3D ROI. This may leave holes in the interior of the ROI if the anchor is not a solid object.
  - Make coincide with intersection: this is interesting to do object analysis with colocalizing volumes only. When you have two pipes active and some objects colocalize (so that the intersection volume of the two pipes is not empty) a ROI can defined with this operation, and applied it to discard objects and parts of objects outside it, so that only the colocalizing regions remain.
  - Envelop anchors: use the currently selected anchors to define a 3D ROI, so that also holes inside the objects are 'filled in' and in the ROI.
  - Enclose anchor in a box: define a prism that is an envelop to the selected anchors.
  - Make a spherical ROI of a given diameter, centered in the image. You can shift it later with the shift ROI mouse mode.
  - Set to all volume: maybe a good starting point for further interactive modifications.
- Modify:
  - Fill inner cavities: a ROI defined by using a threshold may contain inner cavities (visible when selecting the deep Transparency Depth). This operation fills them in automatically.
  - Fill inner and cutoff cavities: A cutoff cavity is a hole in the surface of a ROI that touches the image limits.
  - Grow/shrink: the currently defined ROI can be enlarged or reduced in a certain number of VoXels, independently in the *xy*-plane or in the *z*-direction. A 3D (*xyz*) reduction is also possible: here the number of voxels in the *xy*-plane will be partially adapted to the entered voxels along *z* to, considering the voxel anisotropy, grow/shrink the volume proportionally.
  - Outer shell: re-define the ROI considering only an outer shell of given thickness.
  - Invert the ROI volume.
- Storage
  - Save current ROI to file
  - Load ROI from file
  - Add ROI from file
  - Intersect with ROI from file
  - Subtract ROI from file
- Center ROI on the anchor CM: align the Center Of Mass (CM) of the currently defined ROI with the CM of the selected anchors.
- Clear ROI
- Keep only objects inside the ROI, discarding anything else. The relaxed selection mode in the options also affects how objects partially inside the ROI are handled.

- Analyze ROI volume computes and reports information on the table about the ROI itself.
- Analyze all objects inside the ROI reports in the table information about objects inside the ROI, or partially outside it, depending on the relaxed selection option.
- Help on Regions of Interest.

ANCHORS

- Select all objects as anchors.
- Deselect all anchors.
- Invert current anchor set.
- Set anchors by filtering. This opens a filter dialog as explained in "Table and Analysis Shortcuts (19)" on page 75, but allowing you to select or deselect anchors instead of discarding objects.
- Keep anchor objects, discarding anything else
- Discard all anchor objects.
- Analyze only objects select as anchors

FILTER: some useful predefined filters to remove objects based on their features, and access to a full control filter tool and to reload the original data.

- Quickly remove objects that are touching the borders of the image, as they are surely incomplete.
- Quickly remove objects based on size or sphericity
- Quickly remove non-colocalizing objects
- Quickly find pairs of objects, inside the same pipe or by combining the two pipes.
- Advanced filtering shows the same filter dialog explained in "Table and Analysis Shortcuts (19)" on page 75.

HELP: shows on-line help and tutorials

Mouse Modes (2)

These buttons control the Mouse mode, what the mouse does when you left-click on the image. When you hover with the mouse over these buttons you get a tip on what they are for. If they have a keyboard shortcut to activate them, it is also shown here. For example, the first mouse mode is intended to *analyze objects*, and you can always activate it from within the rendering view by pressing the keyboard key '1'.

From left to right, and from top to bottom:

- **Analyze object** lets you click on different defined objects and obtain the local statistics.
- **Select area** lets you define a 2D region on the current view of the image so you can do different things with it: analyze or discard objects below it, anchor them as references, or interactively define regions of interest (ROI). We will see what all this is useful for.
- **Discard object** lets you discard irrelevant objects one by one. Just select this mode and click on the disturbing objects.

- **Select object** (as anchor) lets you select and deselect objects to be 'anchors', for example to act as references to measure distances from other objects when asking for local statistics. When you set an object as a reference anchor it will 'light up' and change color on the screen to indicate its new status. It is possible to select a group of anchor objects, and you can operate with them through the Anchors menu.
- **Rotate scene** interacts with the full image to rotate it in the space, by dragging the mouse pointer on the rendering view. That can also be achieved by moving the Tilt and Twist sliders along the rendering.
- **Pan scene** interacts with the full image to move it in space laterally. This means that you can pan the scene in the 2D plane of your screen, not along the third dimension, along your line of sight.
- **Pan canvas** is similar to 'pan scene' but not exactly the same. It allows you to explore the canvas by not re-rendering the scene. This only makes sense when you have a canvas larger than your rendering window, of course. See OPTIONS→VIRTUAL RENDER SIZE in the top menu.
- **Shift the ROI**. This is the only mouse mode that is not always enabled: you need to have defined a region of interest (ROI) before moving it around.

For the advanced users: some of these mouse modes have 'shortcuts' in other mouse modes. Most of the times you can 'pan scene' independently of the selected mouse mode if you use your mouse right button instead of the left one. Similarly, you can 'pan canvas' using the middle button at any time. You will learn this with practice if you need it, but you don't have to care about it right now. When a mouse mode is active and the cursor is inside the canvas, the tooltip in the status bar (See Figure 12.1, item 23) tells you what each mouse button can do.

Selection Interactive Operations (3)

The rest of the buttons in the toolbar are not mouse modes: they do not set new behaviors for your mouse when interacting with the objects view but execute operations on previously defined conditions, for example deleting the objects inside a defined ROI. They are all disabled until these conditions have been set (in the example, until you define a ROI to operate with).

After having drawn a 2D selection (See "The Selected Area (9)" on page 71) in the *Select area* mouse mode, you can click on:

- Keep all objects under the selected area, discarding anything else.
- Discard all objects under the selected area
- Set as anchors all objects under the selected area

The way the drawn 2D area considers objects in its limit can be controlled with an option in the top Options menu: the area can consider only objects fully under the selection, or also objects partially outside it (relaxed selection mode). In any case, the objects not affected by the selection are rendered with dimmer intensity to clearly indicate what objects are selected.

When you have selected certain objects as anchors, you can further operate with them:

- Deselect all anchors.

- Keep the selected anchors, discarding anything else.
- Discard the selected anchors.

More operations are available at the top ANCHORS menu. The main purpose of an anchor is to serve as analysis reference in some experiment presets, for example to measure distances relative to them. But you can also use the selected anchors to delete some objects, or to define a 3D ROI based on them.

ROI Interactive Operations (4)

- Reduce the current ROI to the volume under the selected 2D area (intersects the ROI with the selection)
- Add the volume under the selection to the ROI volume (union of the ROI with the selection).
- Remove the volume under the selection from the defined ROI (difference of the ROI with the selection).
- Set all objects inside the ROI as anchors. This requires a ROI to be defined, of course.
- Keep objects inside the ROI, discarding anything else.
- Discard the defined ROI, reverting to the whole image.

Data Cropping and Restore Operations (5)

- Once a ROI is defined, you can use these scissors to crop the dataset and delete (set to zero) all voxels outside the ROI.
- Export the current data (as cropped by the ROI, or after deleting objects interactively) to the original image in order to save it.
- If deleted data was not yet exported to overwrite the original, you can always reload the original dataset with this button.
- Undo last operation.

Active Pipe Mode (6)

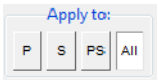

These radio buttons control, for most of the interactive, ROI and selection operations, on which of the active data pipes (*Primary*, *Secondary* or both) and channels (*Selected pipe*, or *All*) they act. Pipes are explained in "Render Pipes (16-17)" on page 74.

Examples of operations controlled by these buttons are:

- Interactively clicking on objects. When you click on the scene, only objects in the active pipe are considered.
- Analyze all / Analyze selection
- Set ROI to the visible objects
- Sum all column values (when hovering over a column header in the table)
- Apply the ROI to crop the dataset

There are only a couple of practical differences between the PS and All modes:

When cropping data by applying a ROI, you can crop the channel in the Primary pipe (P), the channel on the Secondary pipe (S), on both (PS), or all channels in the image (All), even if they are not shown in any active pipe.

When summing cell values in a column of the table, you can include in the summation cells for the channel in the Primary pipe (P), the channel on the Secondary pipe (S), both (PS), or all cells (All), including those reporting about the ROI.

For all the other operations the PS and All modes are equivalent.

Reference Cube (7)

At the end of the buttons column there is a colored reference cube that will help you in orienting in space when you rotate the dataset, specially with large zoom factors that do not let you see the surrounding box frame in the rendered image. The initial view of this cube is the blue top face corresponding to $z = 1$. Hovering the mouse over the cube faces brings a tooltip with the face label (x, y and z with values 0 or 1).

The Interactive Rendering Canvas (8-11)

The canvas shows the scene, the result of the ray tracing[1] algorithm rendering the segmented objects.

The scene is determined by the objects orientation (*tilt* and *twist*), the *zoom*, the *brightness* of the pipes, what point is centered on the view and so on. All that is taken into account by the renderer, that generates the scene, and puts it in the canvas on the screen for you to see it. Notice that depending on the render size (OPTIONS→VIRTUAL RENDER SIZE) the canvas can be larger than your screen.

The On-Screen Reported Parameters and Tooltips (8)

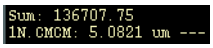

The currently selected Experiment preset (See "Experiment Presets (14)" on page 73) selects the statistics to report many parameters to the table (See Figure 12.1, item 20, 21, and 22) A few of these parameters can be also reported on the screen for the current object for easy reading. The magnitude of the distance that is plotted on the screen is followed by a triple dash '---'

The Selected Area (9)

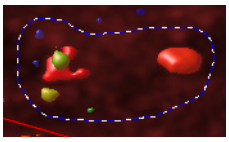

The *Select area* mouse mode lets you define a region in the image so you can do different things with it: analyze objects 'inside' it, anchor them as references, or use the selected volume to define a region of interest (ROI).

Therefore the *selection* is not yet a 3D ROI, but simply a 2D area. That is why 'inside' is quoted in the previous paragraph: objects are inside the selection only from the current point of view, so it would more appropriate to say 'below' it. It allows quick and simple interaction with the objects, specially on flat images here the objects remain more or less

1. http://www.svi.nl/FrontPageRayTracing

in a plane, but only by defining a 3D ROI you have full control on what objects are removed when you want to discard some of them, or when you want to crop the dataset.

In the Select area mouse mode you can use the right mouse button to shift the defined selection around and reuse it multiple times in different locations.

Anchor Objects (10)

Objects that act as references (anchors) are shown in the rendering with magenta or violet colors, depending on the pipe they belong to, so that they are clearly distinct from other regular objects.

There are different ways to set objects as anchors, interactively (See "Selection Interactive Operations (3)" on page 69) or by using filtering operations (top ANCHORS menu). See the on-line article about anchor objects for more details on their utility.

Aim (11)

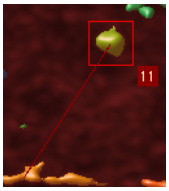

Objects under your cursor are shown highlighted. Objects that have been already analyzed are also shown enclosed by a box when the mouse moves over them, or over the correspondent table row. To analyze an object just click on it while in the *Analyze object* mouse mode, or press the ANALYZE ALL button (See Figure 12.1, item 19).

Objects in the Primary pipe are framed with a red box, and objects in the Secondary pipe with a green one.

A small label showing the object's number ID is also shown when pointing at it. The background color also indicates if it belongs to the primary (red) or the secondary (green) pipe. Yellow labels are shown whenever the pointer has two objects below it, from different pipes.

If a distance is configured to be reported on-screen (See Figure 12.1, item 8), it is also plotted when pointing to an object.

If you hover over an object's row in the table and point to a cell containing a distance parameter, it will also be plotted on the rendering canvas. Like this you can interactively explore many reported distances.

Notice that, depending on the active pipe mode (See "Active Pipe Mode (6)" on page 70), the interaction with the scene may highlight and affect objects in one pipe only.

Scene Control Sliders (12)

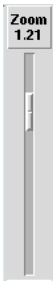

Three sliders run along the canvas (vertically, on its right, and horizontally, on its bottom) to control the point of view of the scene:

- **Zoom**: the number is just indicative, 1 meaning that the whole dataset is shown in the canvas.
- **Tilt**: the angle of rotation (in degrees) around the canvas *x*-axis.
- **Twist**: the angle of rotation (in degrees) around the image *z*-axis.

By clicking on the labels you can enter numerical values manually to quickly switch to the desired scene.

Hide Pane Button (13)

This button at the top right of the window collapses the control pane to make more room for the rendered scene. Once you have defined the segmentation conditions for all pipes, you don't need to interact with those controls anymore, but with the objects, so you can hide the pane to focus on the scene.

Experiment Presets (14)

This button opens a preset selection dialog that allows to select an experiment preset, a collection of parameters that make sense to be reported together in the context of certain experimental needs (See Figure 10.6 on page 56).

Users are very much welcome to send their own suggestions to implement new presets. We will gather all the feedback in different wiki articles that will explain what parameters are best for certain experiments and how to interpret them. Please feel free to tell us about your experiences!

In this dialog you can see three columns. On its left column a series of different experimental needs are listed. When you click on each of them, a different list of parameters is listed in the middle column, and a description is shown on the right column. Even more, when you hover with your mouse over the listed parameters you get a tooltip text explaining each parameter with more detail.

By selecting a particular preset all the listed parameters (apart from the basic ones) will be reported and calculated. Moreover, all the listed parameters will also be available for filtering the objects (See "Filtering Objects" on page 75). The check boxes allow a few of these parameters to be also reported on the screen (See Figure 12.1, item 8) for the current object. The magnitude of the distance that is plotted on the screen is followed by a triple dash '---'. Only one distance can be plotted at a time by clicking on an object, but many can be reported on the table and explored interactively there by simply moving the cursor over the cells.

Only one preset can be selected at a time, and all its parameters will be reported. Advanced and more flexible configuration of the parameters is always possible through OPTIONS→CONFIGURE STATISTICS REPORT in the menu. In this configuration dialog users can also store any set of reported parameters as a new preset.

Time frame selector (15)

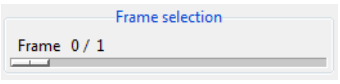

When time series are loaded in the OA, this slider controls which time frame is take for visualization and analysis.

Changes to the current frame (like discarding objects or selecting anchors) are remembered when you change the time frame. But voxel editions (using the ROI to crop the dataset) are lost unless you export them first with *Export the applied ROI to original* ( ).

You can analyze all frames in a time series, accumulating the data in the table, through the FILE→ANALYZE TIME SERIES menu entry.

Render Pipes (16-17)

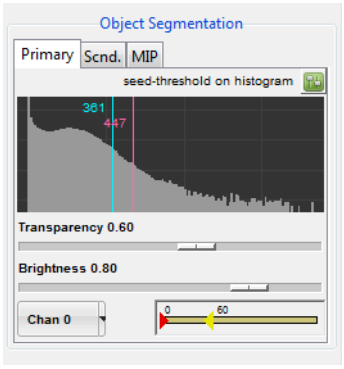

The Huygens Object Analyzer has two pipes (named *Primary* and *Secondary* pipes) for you to put image data through for object analysis and another pipe to simultaneously visualize a MIP of one of the channels (See Figure 12.1, item 17).

'Secondary' does not mean here 'less important', it is just that you can use the first pipe (*Primary*) and the second pipe (*Secondary*), abbreviated *P* and *S* in some places. We keep numbers to refer to image channels, which is something different, because in each of these pipes one can put any image channel, in any order. One can even segment the same channel twice with different conditions in each of the pipes if necessary!

This pane controls what data channel goes through each pipe, how its objects are segmented (with a seed and a threshold[2]), and how the data is rendered (transparency and brightness). A garbage volume in voxels can also be set to get rid of little spurious objects, with a number of voxels lower than the garbage level.

You can see the seed as a secondary threshold. The first threshold segments the data and makes independent objects, but then only objects with intensity that goes above the seed level remain, while the rest are discarded.

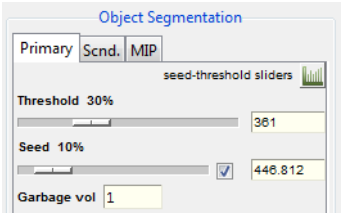

At start-up, the threshold and seed levels are represented by vertical blue and magenta lines, respectively, on top of a histogram of the channel in the pipe. An alternative representation of these levels as sliders is available, that also show their values as a percentage of the channel maximum and allows you to link the seed with the threshold. The garbage volume control is also in the slider pane, not in the histogram. To switch between the histogram and the sliders control panes you have to click on the small button ( ) at the right of the pane title.

The threshold ranges between the minimum and the maximum values in the channel intensities, considering all the time frames. Its percentage representation refers to the maximum value.

By default, the seed is linked to the threshold value, so when you vary the latter the seed absolute value also changes, in such a way that its relative value remains constant. In its linked mode, the seed is set referred to the span between the actual threshold and the maximum, and ranges from the threshold value itself (0%) to the image maximum (100%). This is because the seed, being an 'upper threshold', can never be lower that the threshold. It is also useful to be able to set the seed to 0% so that it is not used at all and you retrieve the classical threshold-only segmentation. Still you may find convenient to express the seed relative value in the same terms you use for the threshold (relative to the image intensity rage), or to be able to set the seed to a fixed value independent of the threshold. For that, deselect the seed checkbox so that 0% also represents the same value as the 0% threshold (the image minimum, or zero), and the two sliders are unlinked. The

2. `http://www.svi.nl/FrontPageSeedAndThreshold`

linking mode of the seed does not really affect the segmentation, it is just a matter of convenient representation of relative values: what is applied to the image as threshold and seed are always the absolute numeric values shown on the entry widgets and next to the histogram lines, that you can edit directly by clicking on them.

All the different objects in a pipe are colored differently to be able to distinguish them. The range of colors assigned to each pipe can be controlled with a *Hue Selector* (See "Hue Selector" on page 105). You can collapse this range completely if you want so all objects in a pipe get the same color.

General Object Information (18)

This little bar reports:

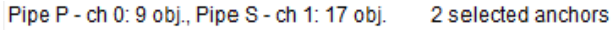

- left: the number of objects currently segmented in each surface pipe.
- right: the total number of objects selected as anchors.

Table and Analysis Shortcuts (19)

These widgets give quick access to some table operations.

- Auto clean checkbox: this option makes the table to be cleaned whenever the segmentation or analysis conditions change, to always have a fresh start. Deselect this option to keep all data and keep accumulating rows in the table.
- Clear statistics table: manually delete the table contents
- Filter: opens a dialog that allows you to filter objects out based on the reported parameters (see below).
- Analyze all / Analyze sel.: This button runs the analysis procedure on all the currently segmented objects, for objects on the pipes selected by the Pipe mode radio-buttons (See Figure 12.1, item 6). When a 2D area has been selected (See Figure 12.1, item 9) this buttons analyzes only the objects under the selection.

Filtering Objects

The Filter button opens a pop up dialog that allows you to discard objects:

- Based on any of the reported parameters...
- Using a certain arithmetic operator...
- To compare the parameter with either a fixed value or with another reported parameter

The pipes that are filtered can be controlled with the active pipe radio-buttons.

You have to select one of the two options (fixed value or another parameter) and enter the value or select the parameter you want to compare with.

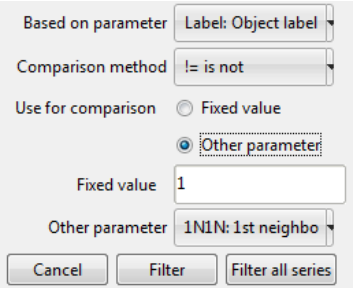

Because the filter is based on the reported parameters, you may need to configure the reported parameters or select other statistics first in order to filter based on the desired property.

For interesting usages of this filter, see for example the neighbors article in the SVI-wiki[3].

For time series, another button shows to allow filtering all time frames.

Statistics Table (20-22)

The statistics table is the place where all the objects parameters are reported after the analysis. You can explore the table values in interesting ways directly on the Object Analyzer (see below), or export the contents for further analysis in other program.

When you move your mouse over the table rows, the objects they refer to are highlighted on the rendering canvas. When you point to a cell reporting a distance, this distance is also plotted on the screen.

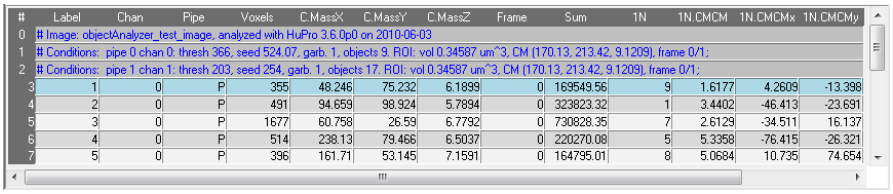

#	Label	Chan	Pipe	Voxels	C.MassX	C.MassY	C.MassZ	Frame	Sum	1N	1N:CMCM	1N.CMCMx	1N:CMCMy
0	# Image: objectAnalyzer_test_image, analyzed with HuPro 3.6.0p0 on 2010-06-03												
1	# Conditions: pipe 0 chan 0: thresh 366, seed 524.07, garb. 1, objects 9. ROI: vol 0.34587 um^3, CM (170.13, 213.42, 9.1209), frame 0/1;												
2	# Conditions: pipe 1 chan 1: thresh 203, seed 254, garb. 1, objects 17. ROI: vol 0.34587 um^3, CM (170.13, 213.42, 9.1209), frame 0/1;												
3	1	0	P	355	48.246	75.232	6.1899	0	169549.56	9	1.6177	4.2609	-13.398
4	2	0	P	491	94.659	98.924	5.7894	0	323823.32	1	3.4402	-46.413	-23.691
5	3	0	P	1677	60.758	26.59	6.7792	0	730828.35	7	2.6129	-34.511	16.137
6	4	0	P	514	238.13	79.466	6.5037	0	220270.08	5	5.3358	-76.415	-26.321
7	5	0	P	396	161.71	53.145	7.1591	0	164795.01	8	5.0684	10.735	74.654

The contents of the table can be copied to your clipboard, or stored to a file in disk by using the FILE menu.

The Table Columns and Their Headers (20)

Distance to first neighbor, same pipe (um)

When you move your mouse over the column titles at the top of the table you get a description of each parameter, at the very bottom of the window (See Figure 12.1, item 20). There are many parameters that measure complex things, so in the column title there is only room for a cryptic label. You can always find out what each parameter is by looking at its tooltip. The description of each parameter will be also stored in your file when you export the table later.

This brief parameter description, plus longer explanations, are also given in the selection and configuration dialogs that selects or configures the reported parameters.

For columns with numeric values, basic descriptive statistics of all values in the column are also reported along with the parameter description. You can select whether the sum runs for both pipes (ALL or PS) or only for one of the pipes (P or S) depending on the selected pipe mode (See Figure 12.1, item 6). This provides a quick way of finding the total number of voxels in the object list, or those that are colocalizing, for example. Many interesting questions can be answered by this summations and ratios between them.

You can right click on a column to pop-up a contextual menu, from which you can retrieve more detailed statistics (STATS) for that column. The descriptive statistics will analyze pipes separately and together, and also include ROI information if present. The reported values are the maximum (MAX), the median (MED), the minimum (MIN), the number of items considered (N), the summation of the values (SUM), the average value

3. `http://www.svi.nl/FrontPageObjectAnalyzerNeighbors`

(AVG) and the standard deviation for the *N* items (SD-N), reported also as a percentage relative to the average value.

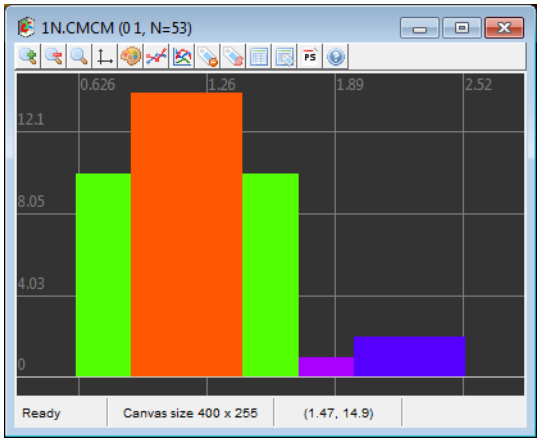

Clicking on a column title selects the whole column for you to copy. Multiple columns can be selected by holding the `Ctrl` key. To plot a histogram of the distribution of values in a column, select the column (or a subset) and select Histogram from the pop-up menu.

Selected columns also act as 'special parameter selector'. The value of the parameter for selected columns will be reported on the status bar (See Figure 12.1, item 23) when you explore cell values for given objects. Like this, you can quickly compare different parameter values for the same object, by looking at the report in the status bar while you move your mouse over the table cells.

A similar thing can be also done by selecting rows, see below.

When multiple columns are selected, the STATS popup dialog will report descriptive statistics for all them.

Conditions Report (21)

When new segmentation or analysis conditions are set, they are reported when you analyze one or all objects. One row is added to describe the image and the time of the analysis, and another row per active pipe is added that reports the image channel in that pipe, the segmentation conditions, the number of objects in that pipe, and the volume and geometrical center of mass of the ROI:

```
# Conditions:  pipe 0 chan 0: thresh 504.8, seed 540.52,
garb. 1, objects 12. ROI: vol 1103 um^3, CM (130, 115,
10)
```

The parameters about the ROI (volume and CM) are the same for both pipes, in case two are active. This is because there is only one ROI for all pipes and channels, and the Center Of Mass (CM) here reported is calculated considering the ROI as a uniform, solid object.

There is another way to calculate the CM of the ROI, not considering it homogeneous but taking into account the real image intensities in envelops. These intensities, being different per pipe, yield to different ROI's CM per pipe. This is not reported in the conditions rows, but as object rows in the table when you select Analyze all. In this case, the ROI itself is treated as a new object in each pipe and more detailed information is reported in separate rows in the table. See the wiki article about ROI[4].

4. `http://www.svi.nl/FrontPageObjectAnalyzerROI`

The Table Rows, How to Explore Them (22)

Each row in the table is either a *Conditions report* (See "Conditions Report (21)" on page 77) or information about one of the segmented objects.

In the second case, the object the row refers to is identified by some mandatory parameters (parameters that are always reported):

- *Label*: an integer index that is unique, per object, inside its pipe.
- *Chan*: the image channel this object belongs to
- *Surf*: the surface pipe in the current analysis the object belongs to (*P* is *Primary*, *S* is *Secondary*).
- *Voxels*: the number of voxels in that object
- *C.Mass*: three columns (*x*, *y*, *z*) for the center of mass location of the object, in the image coordinate system, with voxels as units.

The second letter in the *Surf*. column informs whether the object you clicked was in the front (*F*) or the back (*B*) position. This is only relevant in the following situation: it may happen that two objects in different pipes overlap in space (or they apparently do from the current point of view) and when you click somewhere on the screen you are actually selecting both of them, if the pipe interaction mode is PS or ALL (See "Active Pipe Mode (6)" on page 70). In that case, two rows are added to the table, and this second letter lets you know which of the two was in front of the other, from the current point of view.

By moving the mouse over the table rows the corresponding object is highlighted on the canvas (See Figure 12.1, item 11). Objects in the Primary pipe are highlighted in red, and objects in the Secondary pipe in green.

A good way to find an object in a very long table is by clicking on it on the rendering canvas while the *Analyze object* mode is active: the table will be shifted to show its corresponding row, and it will be highlighted.

When you move over the table cells the current value is shown on the tooltip bar at the bottom (See Figure 12.1, item 23). This, combined with the selection of rows (see below) or columns (See Figure 12.1, item 20), allows a quick exploration of the table and the comparison of different parameter values.

In the following example, when the cursor is moved over one cell that contains the distance to the first neighbor, the following is reported in the tooltip: `38/S 1NP.CMCM: 4.2358 um`. The first part is the label of the object in the current row: label 38 on the Secondary pipe. Then the parameter tag 1NP.CMCM refers to the CM-to-CM distance between this object and the nearest object in the other pipe. Then comes the distance itself: 4.23 microns.

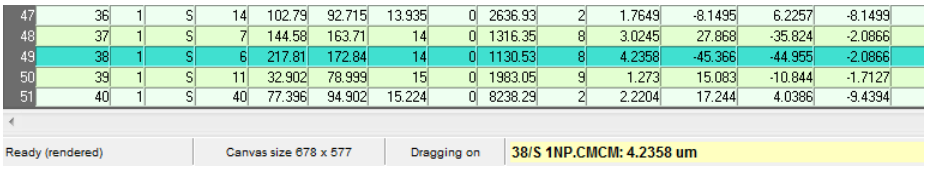

That tooltip region can show not only the information of the cell pointed by the mouse, which would not be really useful, but also other values that can be set as reference. Try this: while keeping the `Ctrl` key pressed on the keyboard, select a couple of rows by clicking on the row number at the very left of the table. The selected rows will turn

green. If you now hover the mouse over a cell on any other row, you will get the value not only of that cell, but also those in the corresponding cells of the selected objects (rows). This is a quick way to compare results for different objects that can be distant in the table.

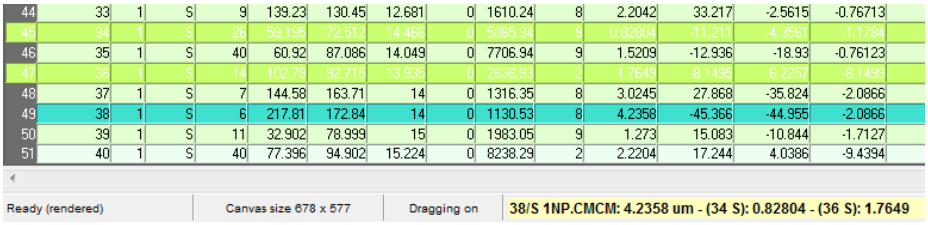

If you select columns instead of rows in the table the tooltip will display the corresponding parameter values for the same object. Do not highlight columns and rows at the same time by now: it may be too confusing to interpret!

The Status Bar and Tooltip (23)

The bottom part of the window is a status bar that also shows an contextual tooltip.

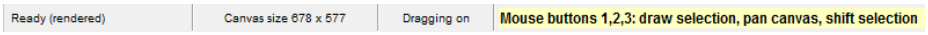

The left side reports the current status of the renderer and analyzer. You can see whether a long computation is running or if the analyzer is ready for further interaction.

Then the current size of the canvas is reported. When you first start the OA, the canvas size is adapted to the exact room left for the scene rendering, but it can be larger or smaller at wish (OPTIONS→VIRTUAL RENDER SIZE).

The 'Dragging' status refers to the automatic rendering of the scene while the user interacts with it. In very slow systems, the dragging is turned off automatically and the rendering only happens after the user released the mouse buttons or finished moving the segmentation sliders.

The tooltip (the region with light yellow background) shows contextual information:

- A longer description of the reported parameters when you hover over the table headers (See Figure 12.1, item 20) and a sum of the cell below them.
- The value in the cell under the cursor when you point at table rows, plus selected reference values.
- The different actions bound to the mouse buttons, when the cursor is inside the rendering canvas.

CHAPTER 13

The Colocalization Analyzer

The *Colocalization Analyzer* provides information about the amount of spatial overlap between structures in different data channels, for 3D images and time series.

As this overlapping can be defined in many ways, Huygens gives the *colocalization coefficients* most commonly used in literature: (*Object*) *Pearson, Spearman*, *Overlap*, *Intersection*, and *Manders M* and *K*. More information on these coefficients can be found in the SVI support Wiki[1,2].

The *Colocalization Analyzer* is an extended optional tool, and is enabled by a C flag in the license string (See "License String Details" on page 101).

Notice that the Object Analyzer (See page 49) also provides colocalization measurements at the object level. The Colocalization Analyzer works more at the level of the whole image, despite local statistics of the colocalizing regions can be easily retrieved. Both analyzers work, in a sense, in complementary ways.

The Object Analyzer allows to define objects (segmentation) and observe how much they overlap, in volume or intensity. Objects defined like this can overlap with other objects, or not.

The Colocalization Analyzer explores the whole image to search for colocalizing regions based on the usual colocalization coefficients. These regions are then segmented and treated as objects to analyze. These objects are therefore always volumes of intersection.

How to use the Colocalization Analyzer

To *start* the Colocalization Analyzer in Huygens Essential, right-click on an image's thumbnail to open the contextual menu, then select COLOCALIZATION ANALYZER. Alternatively, an image's thumbnail can be selected, then in the menu bar select VISUALIZATION → COLOCALIZATION ANALYZER. The image must be multi channel (See "Multi-channel Images" on page 16) as the colocalization is based on the overlapping of different channel intensities. Wait for the analyzer to initialize and to compute the first 2D histogram with the default settings.

1. http://www.svi.nl/FrontPageColocalizationTheory
2. http://www.svi.nl/FrontPageColocalizationCoefficients

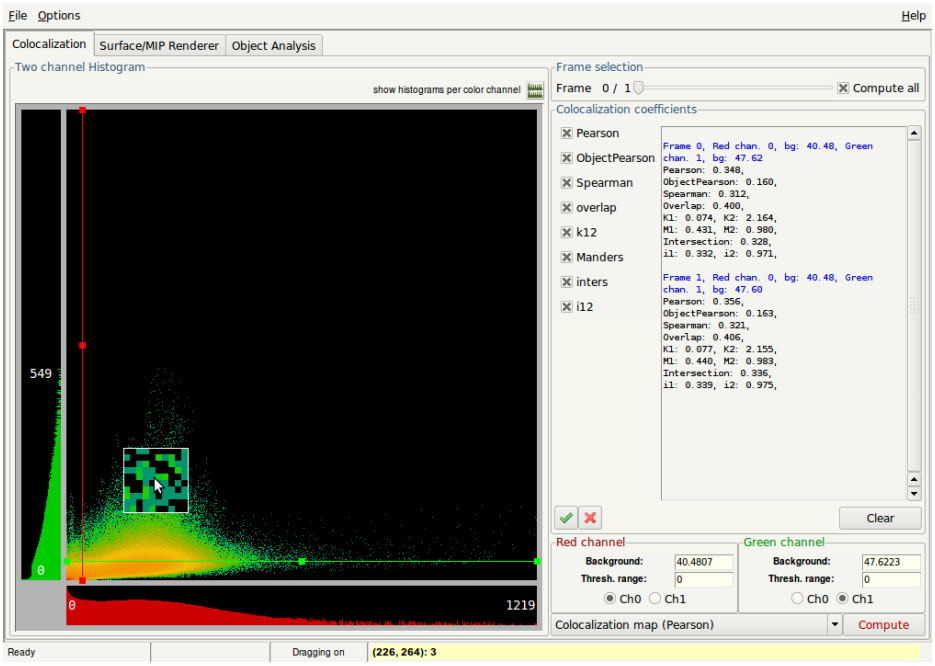

FIGURE 13.1. The first tab of the Colocalization Analyzer showing the colocalization coefficients and histograms.

When the colocalization analyzer is opened you will see three tabs (See Figure 13.1), of which the first is visibile. In this tab the colocalization coefficients can be calculated. You can choose the timeframe, select colocalization coefficients, set the background settings and choose a colocalization map. In default start-up, all these settings are set for you, such that you can inmediatly press COMPUTE.

First we select the data to analyze. For time series, the *Frame* slider selects the time coordinate.

We follow the usual naming convention in colocalization theory for the two compared channels: *Red* (R) for the first channel, *Green* (G) for the second channel. We can select, in the lower part of the window, which data channels from our image are the Red and Green channels to be compared.

A two-channel *histogram* is calculated by default, and updated whenever we change the Red or Green channel selection. This histogram is already an indication of the degree of overlapping between the selected channels: for two channels with a high degree of overlapping, the histogram pixels trend to concentrate along the diagonal $y = x$ line. In contrast, total absence of overlapping would produce a 2D histogram with values only on the coordinate axes1.

Along the x and y line of the 2D histogram, there are 1D histograms for each color channel separately. The enlarged versions of these 1D histograms are shown, when clicking on the SHOW HISTOGRAM PER COLOR CHANNEL button at the top right of the histogram. By moving the colored *background* lines in either the 2D-histogram, the enlarged 1D-histograms, or by changing the numeric values in the input fields, you can specify the backgrounds for the two selected channels. These values are subtracted from the voxels intensities when calculating the coefficients (if the result is negative, it is understood as a zero). Generally the colocalization coefficients depend much on correct estimation of

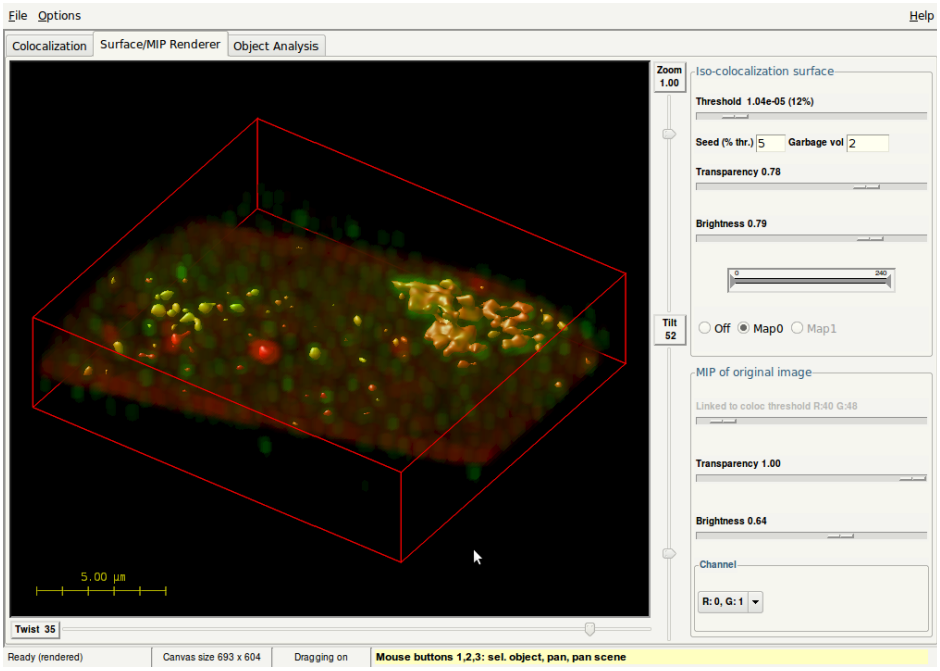

FIGURE 13.2. Tab two of the Colocalization Analyzer shows surfaces for regions of overlapping channels.

the image background and resolution. For these reasons we strongly recommend to compute colocalization coefficients only on deconvolved images.

One sets a background value prior to the calculation to remove signal that would otherwise lead to spurious colocalization, . This is intended for minor tuning, or just in case it is really needed to calculate colocalization in raw images that still have all the measured background. See “Backgrounds vs. thresholds in colocalization” on page 84.

Next we select what *colocalization map* we want to calculate: (*Object*) *Pearson, Spearman*, *Overlap*, *Manders M* or *K* or *Intersection*.

Notice the difference between *maps* and *coefficients:* the *colocalization coefficients* parametrize the degree of colocalization of the full image, while a *colocalization map* parametrizes the colocalization locally. In a map, a single colocalization value is calculated per voxel creating a 3D distribution that is represented in a 3D image by iso-surfaces. In other words, the colocalization map shows the contribution of each voxel to the colocalization coefficient.

The Colocalization Analyzer computes only the map selected by the user, but it always computes the selected coefficients. When the COMPUTE button is pressed, the pane *colocalization coefficients* will show the selected colocalization coefficients for the selected time frame. If you have checked the COMPUTE ALL option next to the timeframe slider, the coefficients will be calcuted for all time frames sequentially.

Your choice of the colocalization map is needed for the functionality in the second tab (See Figure 13.2). In this tab an iso-colocalization surface is generated based on the colocalization map. This iso-surface rendering is only possible if a colocalization map exists. So, at startup the iso-colocalization surface sliders are deactivated, as we have to calculate a colocalization map first, by pressing COMPUTE.

Together with a surface renderer a MIP renderer is available. The viewpoint of the MIP and surface can be selected by moving the Tilt and Twist slide. Also try changing the zoom.

The obtained colocalization map is represented in the renderer window by iso-colocalization surfaces. These iso-surfaces represents points which all have the same colocalization value, thus regions in which the degree of colocalization exceeds a certain value become objects. This "certain value" can be controlled by the *threshold* slider in the iso-colocalization surface parameters. The transparency and the brightness of this surface pipe can be controlled with the correspondent sliders. The color range in which the objects are displayed can be modified using a hue selector (page 105). One can also switch the surface pipe off.

Some modes generate two-channel *colocalization maps*: colocalization of red with respect to green, and vice versa, e.g. in case of the *Manders M1* and *M2* coefficients. In these cases, the iso-colocalization surface parameters will offer the possibility of rendering any of the two channels, and thus the threshold is referred to the active one.

By clicking on the rendered objects local colocalization parameters are computed and reported, which will be discussed in more detail in "Iso-colocalization object analysis" below.

There is one *Maximum Intensity Projection* (MIP) pipe available to redirect the data channels to. The *MIP rendering* of one channel (maybe one different from those used for colocalization) or the two channels under analysis can be a good *spatial reference* for the objects from the colocalization map. When an original channel is selected, the *threshold slider* can be used to select what voxels are considered for the MIP rendering, depending on their intensities. Notice that here the threshold is simply used for *representation*. If both R and G channels are selected to be rendered, their correspondent backgrounds as selected in the histogram will be used as projection thresholds. As with the surface pipes, the *transparency* and the *brightness* of this MIP can also be controlled.

All the obtained information can be saved to external data (text or image) files through the FILE menu.

Iso-colocalization object analysis

One of the features of the colocalization analyzer is iso-colocalization *object analysis* (Figure 13.3 on page 85). It allows to determine quickly the properties of the different colocalization regions in the data. This is realized by visualizing the colocalization map as iso-colocalization surfaces. In this way regions in which the degree of colocalization exceeds a certain value become objects. By clicking on the objects local colocalization parameters are computed and reported. To relate the iso-colocalization objects to the original data the surface objects can be blended with a MIP projection of the data (See "The MIP Renderer" on page 29). The color range in which these objects are displayed can be modified using a hue selector (page 105).

Backgrounds vs. thresholds in colocalization

Backgrounds are for removing signal prior to the calculation. In an ideally restored image that would not be necessary, because all the signal present in the image gives a valid representation of it. The background would have been removed during the image restoration.

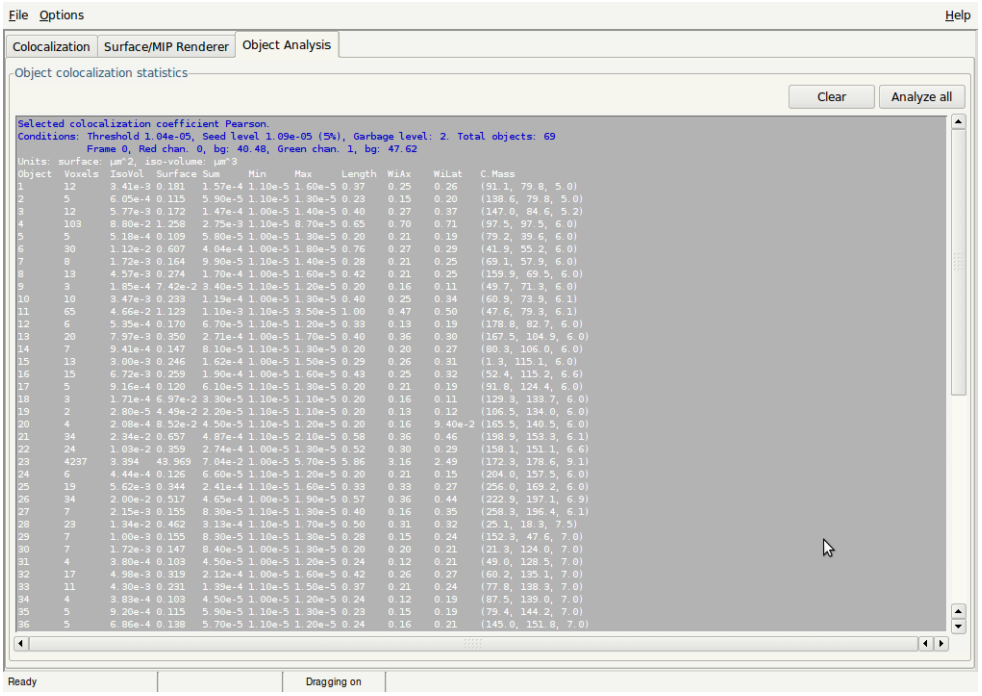

FIGURE 13.3. Tab three of the Colocalization Analyzer if for object analysis for the colocalized regions of the colocalization map.

Thresholds are used to split the colocalization maps in two regions: what are interesting objects, and what are not. Local colocalization values are calculated for every image voxel, but the zero value would be very rarely achieved. In most of the cases there are some non-zero colocalization everywhere, the interest probably is in regions where colocalization exceeds typically a value of 20 % of the maximum.

With the background settings, voxels are in- or excluded based on their voxel intensities: intensities higher then the background are kept, while lower intensities are disregarded for the calculation. Note, the intensity value of voxels do not relate to the level of colocalization. In fact, it may well be that high colocalization levels occur in regions with low voxel intensities.

Therefore, background settings should not be used to remove signal that can still have some colocalization level. Ideally work with deconvolved images and consider all the signal, then study the colocalization levels using thresholds to split regions of high colocalization from regions with low colocalization, something that can not be know beforehand! Setting the background is used for removing really constant background signals.

Thresholds do not affect colocalization, but only the way colocalization maps are represented on the screen and objects are defined. Depending on the locally calculated colocalization objects may merge or split. What affects colocalization, as explained in the colocalization theory at the SVI Wiki, are the backgrounds. In the computation of *Manders* coefficients the background values act like in this example: to the computation of *M1* only pixels in *R* contribute when their corresponding pixel in *G* is above the background.

More information can be found at the SVI Wiki[3].

3. http://www.svi.nl/FrontPageColocalizationAnalyzer

CHAPTER 14

The Chromatic Shift Corrector

The Chromatic Shift Corrector is a post-deconvolution tool that can estimate and correct for chromatic shifts, removing the existing misalignments across different channels. The result of this correction is a channel-aligned image free of chromatic shifts.

The support for templates included in this tool allows to apply the chromatic shift correction of one image to other images. This is particularly interesting when the estimation carried out on a beads image is suitable for other sets of images.

The Chromatic Shift Corrector shows its Help content dynamically to guide the user through the process of estimating and correcting for chromatic shifts.

Starting the Chromatic Shift Corrector

- Launch Huygens Essential or Huygens Professional.
- Load a multichannel image to be corrected for chromatic shifts.
- Select the image thumbnail and in the top menu go to DECONVOLUTION->CHROMATIC SHIFT CORRECTOR.

The Chromatic Shift Corrector will open and show the image on an orthogonal slicer where the existing chromatic shifts can be seen in a 3D view. Below the orthogonal slicer a Z slicer, a time slicer, and other visualization tools such as contrast, colour scheme, channel selection and zoom tools can be found. These tools are useful to enhance the view of the image for a better visualization of the chromatic shifts.

A view of the Chromatic Shift Corrector at start-up with a loaded two-channel bead image is shown in Figure 14.1.

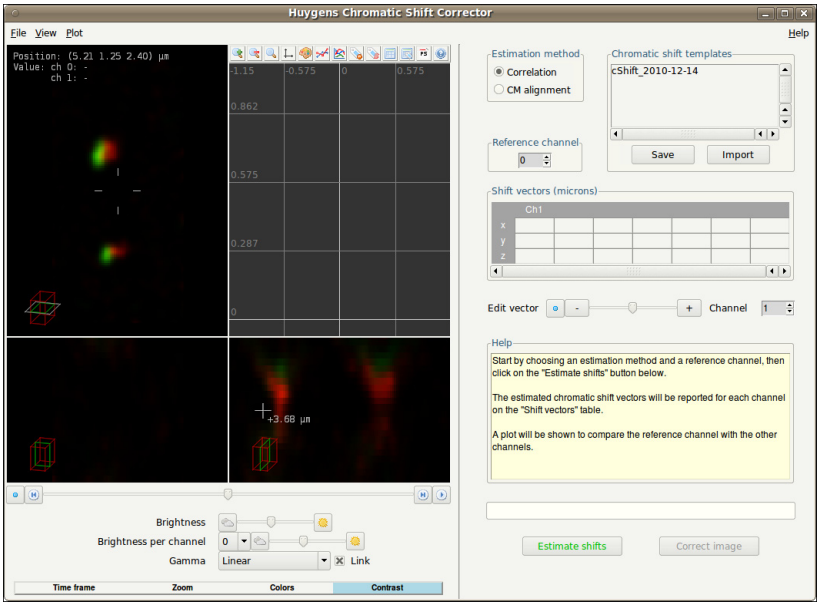

FIGURE 14.1. The Chromatic Shift Corrector at start-up, prior to estimating the chromatic shifts.

Estimation of the chromatic shifts

For the automatic estimation of the chromatic shifts the following two methods can be chosen:

- Cross correlation. This can be considered an 'all-round' method. The software searches for the best alignment across channels by maximizing the overlap.
- Center of mass alignment. This method works best if the image contains a single object. The object should not touch the image borders, and the contrast between object and background should be high.

The chromatic shifts will be quantified by vectors, indicating how much a channel is shifted with respect to a given reference channel.

The channel to act as reference (no shift) can be selected by the user via the Reference Channel selection box. Because this channel will have no chromatic shift its shift will not be reported.

If an estimation method and a reference channel have been selected, the chromatic shifts will be estimated and reported upon clicking on the ESTIMATE SHIFTS button.

Visualization of the chromatic shifts

The estimated chromatic shifts will be reported on the shift vectors table. The shifts of each channel but the reference will be listed. The length unit of the shift vectors is set to micrometers.

The user can select any channel but the reference for editing. This can be done by using the EDIT CHANNEL selection box under the vector table.

A vector will be drawn on the orthogonal slicer showing the estimated shift vector of the edited channel, so that the estimated shift can be easily assessed. At the same time, a plot

will be displayed showing the intensity profiles of the edited channel and the reference channel along the direction of the estimated shift vector.

The plot also shows the intensity profile of the edited channel as if it were corrected with the existing estimated shift (dashed line).

Therefore, the plot serves as a comparison between the reference channel, the edited channel and the corrected edited channel.

This comparison allows us to see to what extent the intensity of edited channel is shifted with respect to the reference channel. At the same time, it shows in advance whether the estimated shift vectors will correct for the existing chromatic shifts properly.

The result of a chromatic shift estimation is shown in Figure 14.2. The shift between the intensity profiles of the reference channel and the edited channel is visible in the embedded plot. Additionally, the dashed line in the plot shows the intensity profile of the edited channel as if it were corrected with the estimated shift.

Ideally, the intensities of the corrected channel (dashed line) and the reference channel would show no gap or shift, having similar shapes and peaks roughly located at the same x positions.

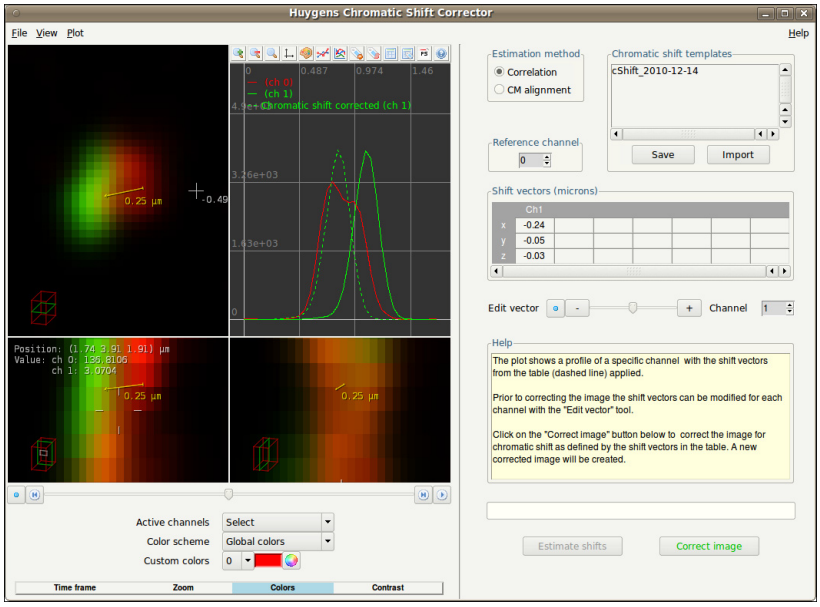

FIGURE 14.2. The Chromatic Shift Corrector after estimating the shifts. A plot shows the gap between two channels. The estimated shift is drawn as a vector on the slicer.

The Chromatic Shift Corrector will return accurate and reliable estimations of the existing chromatic shifts. Nevertheless, the possibility to edit and customize the estimated shifts exists so that the user can reach more precission if necessary.

Editing the chromatic shift vectors

The user can also edit the shifts estimated automatically by the Chromatic Shift Corrector. To apply a different customized correction a channel has to be selected for edition. The components of the selected shift vector can be modified by using the edit tool, which allows to shorten and lengthen the estimated shift.

While the shift of the edited channel is shortened or lengthened the plot updates itself to show how the gap between the reference channel and the edited channel is increased or decreased.

In this way, the contents of the estimated shifts can be modified while checking in advance whether the image will be corrected appropriately.

The chromatic shift estimation of a four-channel image is shown in Figure 14.3. The channels can be edited one by one and their shifts modified with the EDIT CHANNEL tool.

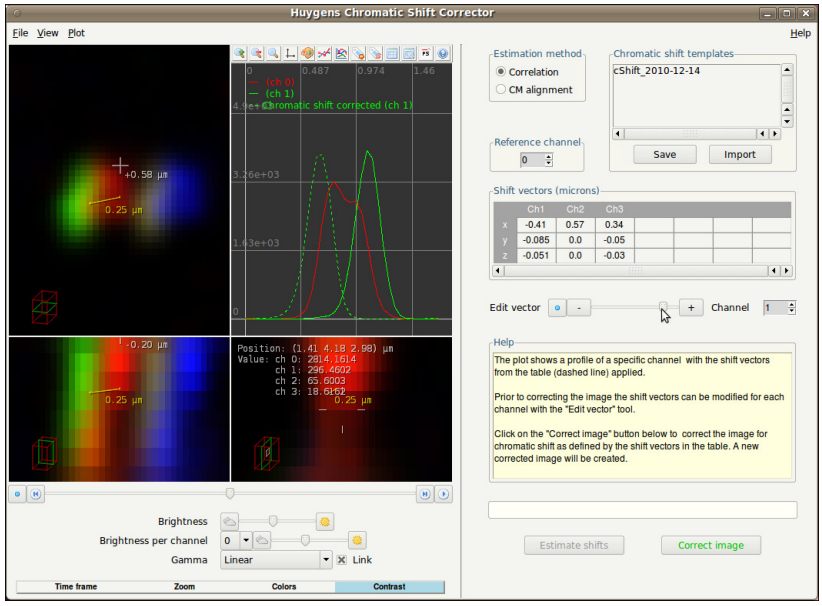

FIGURE 14.3. Using the Edit Channel tool of the Chromatic Shift Corrector. The estimated vector has been lengthened, and the dashed line plot is consequently slightly moved to the left.

Alternatively, any customized shifts can be typed in the vector table for each channel.

Upon clicking on the CORRECT IMAGE button the image will be corrected with the chromatic shifts listed in the vector table. A new corrected image will be created.

Working with templates

The estimated shift vectors can be saved to a template by using the template tool of the Chromatic Shift Corrector.

The saved templates can be imported again to apply its shift vectors to other images. The template vectors will be loaded and listed in the vector table, showing the intensity profiles as if the ESTIMATE SHIFTS button had been pressed.

The image can then be corrected by clicking on the CORRECT IMAGE button, which will create a new corrected image.

In order to be able to work with the Chromatic Shift templates properly it is recommended to apply these templates to images that have the same emission and excitation wavelengths as the image of the template.

CHAPTER 15

Huygens Remote Manager

In this chapter the Huygens Remote Manager (HRM) is presented. HRM is a collaborative open-source interface to Huygens Core that allows scheduled multiuser deconvolution through a web server.

To run deconvolution jobs, HRM will request from the user a number of microscope and restoration parameters. For more specific information on these parameters, visit the online SVI-wiki [1] or follow the HRM help.

This chapter is an overview of the interface possibilities of HRM. Being a collaborative project, the HRM capabilities may change and expand very quickly. It is advisable to follow the HRM online help.

What is HRM

The Huygens Remote Manager (HRM) is a web task manager for servers that acts as an interface to Huygens Core to do multiuser batch deconvolution.

Huygens Essential and Huygens Professional have their own integrated scheduler, the Batch Processor, which is more intended for single-user deconvolution. Multiple users may run simultaneous sessions of the Batch Processor but the multiple processors will compete for the same hardware resources, likely resulting in a slowdown.

HRM, however, has a queuing system intended for multiple users. Each user has his own account in a web server and can place deconvolution jobs in the queue. HRM runs all jobs listed in the queue setting priorities across them and alternating over all users. A quota system for improved queue management can easily be implemented, and will probably be included in future versions of HRM.

HRM is flexible enough to control different computation servers and split the jobs among them, allowing centralized administration of the deconvolution jobs in a cluster. The HRM queue manager is a so-called 'daemon' that runs in background in Unix-like systems (Linux and Mac OS X). HRM is not the only way to use Huygens Core, which is

1. http://www.svi.nl/FrontPageMicroscopicParameters
   http://www.svi.nl/FrontPageRestorationParameters

available for Linux, Mac OS X and Windows. Because HRM is an open source project, it can be freely modified and its code reused to adapt it to any particular needs. Other completely different interfaces that communicate with Huygens Core to use the Huygens restoration, visualization and analysis algorithms could be developed using the Huygens Core Programmer Guide [2].

Huygens Core works by default without graphical interface and is designed to work seamlessly with HRM. It will include more and more features especially intended for web interfaces such as HRM, or any other alternative interface.

Where to find HRM

The open source HRM is developed by Huygens users at the Montpellier Rio Imaging facility, the Facility for Advanced Imaging and Microscopy at the Friedrich Miescher Institute (FMI, Basel), and the BioImaging and Optics Platform at the Ecole Polytechnique Fédérale de Lausanne. Scientific Volume Imaging participates in this project by contributing its experience in deconvolution and software engineering. HRM is a free and open source project, and can be found in SourceForge [3].

More information about HRM and links to other HRM resources can be found in the HRM online article in the SVI-wiki[4]. Instructions for online testing, downloading and installing the HRM code are also linked on that page.

On a running regular web server (including sendmail, Apache, PHP and a relational database like PostgresQL or MySQL) installing HRM is not very difficult. Apart from the installation instructions that come along with the source code, other practical guidelines based on other users' experiences can be found in the SVI-wiki [5].

HRM requirements and technical features

HRM consists of two main components, both written in PHP: a web based interface and a queue manager. The web interface allows:

- the management of different users by the system administrator;
- the management of template settings that all users can copy or use directly;
- the creation of deconvolution jobs by the users, including image selection, setting of microscope parameters, and setting of restoration parameters;
- inspecting the job queue status, and deleting the user's own jobs from it.

HRM is equipped with a simple http file uploader/downloader to send the raw images from the user's local machine to the HRM server, as well as to retrieve the deconvolution results from the server. The server administrator can set up the largest file size allowed in these transactions. This kind of customization is typically handled from the PHP initialization files and the HRM configuration.

The jobs created by the users via the web interface are dispatched by the queue manager of HRM, which runs in the backgroundas a daemon, to any of the dedicated servers run-

2. An on-line version is available http://www.svi.nl/HuCoreMan
3. http://sourceforge.net/hrm
4. http://www.svi.nl/FrontPageHuygensRemoteManager
5. http://www.svi.nl/FrontPageHrmInstallation

ning Huygens Core. When the job is finished, an e-mail informs the user that the restored datasets are available.

More specifically, HRM requires a web server with PHP and e-mail capabilities, a database backend to store deconvolution parameters, job descriptions and, optionally, user accounts, a file server to temporarily store input and restored datasets, and one or more processing servers running Huygens Core.

The setup is highly configurable, since (i) the file server, the processing servers and the queue manager can either be all hosted by the same machine or be distributed across two, three or more computers; (ii) two different databases, PostgresQL and MySQL, are supported through the ADOdb Database Abstraction Library [6] and (iii) parts of HRM are optional, as is the case of the embedded user management system.

How HRM communicates with Huygens Core

For each deconvolution task in the job queue the HRM queue manager automatically generates a Tcl script that

- loads the raw image from a source directory,
- applies the microscopic parameters settings to it as defined by the user in a template,
- optionally loads another image containing the microscope Point Spread Function,
- deconvolves the image using the restoration settings chosen by the user,
- stores the resulting restored image in a destination directory,
- and finally writes a tag in that destination directory to let the HRM queue manager know that the job is finished.

This generated script is sent to the Huygens Core for execution by using the command line option `-task`. The HRM configuration files and database store information about where the `hucore` executable is to be found, and the script is passed as an argument to it.

When the job is finished the queue manager optionally sends the user an e-mail announcing the end of the job and its status. This is handled by HRM itself and not by the above mentioned Tcl script.

Multiple jobs can be processed in parallel depending on how HRM is configured, the multiprocessing capabilities of the server and the number of available computation servers.

6. http://adodb.sourceforge.net

Huygens Remote Manager v1.2.3
What's new?
Website
SVI wiki
Help
Welcome
The Huygens Remote Manager is an easy to use interface to the Huygens Software by Scientific Volume Imaging B.V. that allows for multi-user, large-scale deconvolution.
Friedrich Miescher Institute
Facility for Advanced Imaging and Microscopy
Montpellier RIO Imaging
Federal Institute of Technology - Lausanne
BioImaging and Optics platform
Login
If you do not have an account, please register here.
Username: user
Password:
login
created 2004 by Volker Baecker
modified 2006-2010 by Asheesh Gulati, Alessandra Griffa, José Viña & Aaron Ponti

FIGURE 15.1. The login page of the Huygens Remote Manager. Users with an active account can login, new users can apply for an account.

Basic HRM usage

In this section a summary of the basic HRM work flow is described. Help links are accessible in HRM at each stage of the creation of deconvolution jobs.

Registration and login

To become an HRM user one needs to apply for an account in the system. This application can be sent directly via a registration form available in HRM. The application will reach the HRM administrator who will likely grant the user further access to HRM. When filling out the registration form, all entries but the “Request message” field are mandatory. The “Request message” field in the registration form lets the applicant send the HRM administrator a message along with the application.

An e-mail will inform the user as soon as the account has been activated or rejected. Upon account activation, the user can login with the chosen user name and password (see Figure 15.1). For security reasons the password is not shown while the user types. Upon mispelling a name or a password a message will appear stating that the account does not exist. Be aware that the name and the password are case sensitive, i.e. “pierre” and “Pierre” are different names.

User management

This area is enabled for the HRM administrator only. It allows to perform maintenance on the user database. It contains, for instance, a “registrations” area that lists the pending account requests to be either approved or rejected by the administrator (see Figure 15.2).

The existing users are grouped by the initial letter of their user name. This allows the administrator to filter users for further administrative processing, such as edition, deletion, rejection, etc. Alternatively all users can be shown at once.

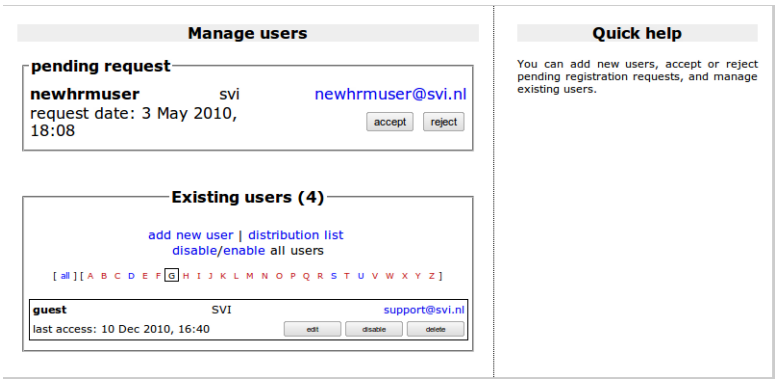
Manage users
pending request
newhrmuser svi newhrmuser@svi.nl
request date: 3 May 2010, 18:08
accept
reject
Existing users (4)
add new user | distribution list
disable/enable all users
[all] [A B C D E F G H I J K L M N O P Q R S T U V W X Y Z]
guest SVI support@svi.nl
last access: 10 Dec 2010, 16:40
edit
disable
delete
Quick help
You can add new users, accept or reject pending registration requests, and manage existing users.

FIGURE 15.2. The HRM User Management Page for the administrator. Active users and new applications can be easily managed from this panel.

New users can be created by the administrator by clicking on "add new user". A user added by the administrator automatically gets access to HRM granted.

Clicking on the "distribution list" link allows the administrator to send an e-mail to all registered users.

Each user can be edited, enabled/disabled, or even deleted by the administrator by clicking on the corresponding buttons next to the user name.

Additionally, the posibility exists to enable or disable all users at once, when seeing the unfiltered list of all users.

Furthermore, the User Management is optional in HRM. In case the administrator chooses authentication via OpenLDAP or MS' Active Directory the User Management is handled outside HRM.

Job settings

The HRM parameter settings are divided into two groups: those describing the image (settings of the microscope) and those describing the restoration process (deconvolution settings). Both are saved in the HRM database.

The HRM settings can be created by the administrator so that they are available across the system, to be used by all users, or can be created individually by each user, either from scratch or based on the settings created by the administrator. The settings created by users are visible only in their account and are not visible across the system by other users. A setting is saved with a name and can be reused later.

Setting both image parameters and restoration parameters is carried out similarly in HRM. For detailed explanations on the meaning of these parameters and how to determine their values please refer to the HRM online help for further details.

Notice again that every stage in HRM comes along with help links with plenty of further information. There are many image and restoration parameters, and understanding them properly is important to achieve good results.

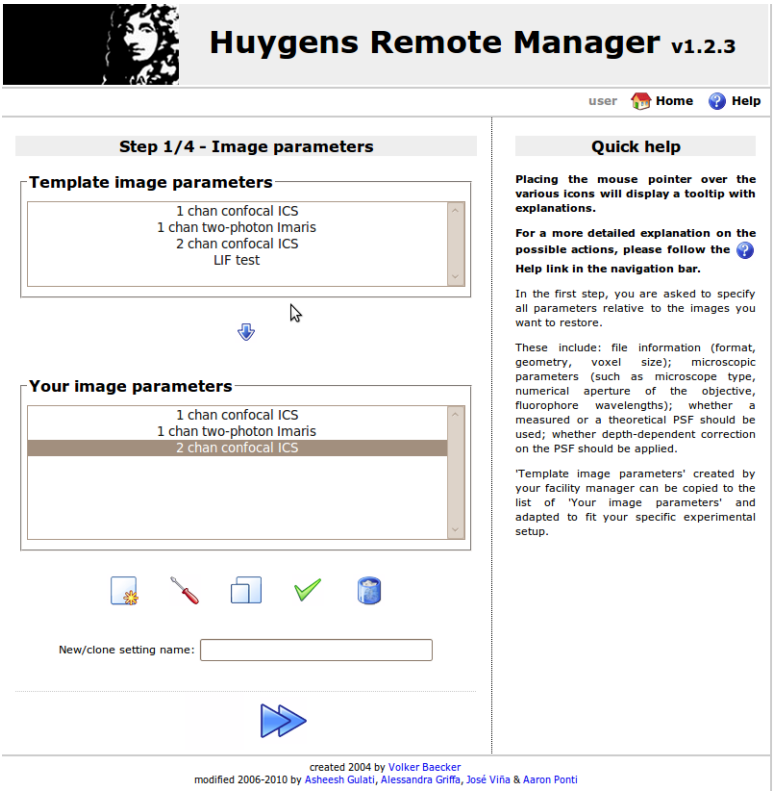
Huygens Remote Manager v1.2.3
user
Home
Help
Step 1/4 - Image parameters
Template image parameters
1 chan confocal ICS
1 chan two-photon Imaris
2 chan confocal ICS
LIF test
Your image parameters
1 chan confocal ICS
1 chan two-photon Imaris
2 chan confocal ICS
New/clone setting name:
Quick help
Placing the mouse pointer over the various icons will display a tooltip with explanations.
For a more detailed explanation on the possible actions, please follow the Help link in the navigation bar.
In the first step, you are asked to specify all parameters relative to the images you want to restore.
These include: file information (format, geometry, voxel size); microscopic parameters (such as microscope type, numerical aperture of the objective, fluorophore wavelengths); whether a measured or a theoretical PSF should be used; whether depth-dependent correction on the PSF should be applied.
'Template image parameters' created by your facility manager can be copied to the list of 'Your image parameters' and adapted to fit your specific experimental setup.
created 2004 by Volker Baecker
modified 2006-2010 by Asheesh Gulati, Alessandra Griffa, José Viña & Aaron Ponti

FIGURE 15.3. Setting selection in HRM. The top pane lists the setting templates created by the administrator for all users. The bottom pane lists settings only available for the current user.

Using an existing setting

Notice that there are two different setting panes on the setting screens (see Figure 15.3). The settings in the bottom pane "your parameter settings" can be modified by the user. The settings template created by the administrator (top pane) can be copied into the bottom pane. Once at the bottom pane the settings can be further worked out to meet the needs of the user. Settings can also be created from scratch if that is more suitable to a particular user.

To copy any of the administrator setting templates from the top pane the template has to be selected and it then has to be transfered to "your parameter settings" (bottom pane) by clicking on the blue down arrow.

The bottom pane "your parameter settings" will list the settings among which the user can choose the most suitable for his or her deconvolution jobs. The selected setting will be highlighted in the list. Press the forward icon to continue to the next step.

Creating a new setting

First enter the name of the new setting. The name must be different from the names of already existing settings. Besides this restriction any name can be used. It is best to choose a name that helps remember what the setting is made for, e.g. it is better to use the name of the used microscope than to use a family name.

Press the CREATE button. Several pages where to enter different sets of parameters will be presented. The number of pages where to set these parameters is 3 or 4, depending on whether corrections for spherical aberration and selection of a measured PSF are chosen. The option to save the setting will eventually show at the last page.

Copying a setting

In order to create a setting that is similar to an already existing setting:

- Select from the bottom pane the setting to be copied. The selected setting will be highlighted.
- Enter a name in the field ”new/clone setting name”.
- Press the COPY button.

The new setting will be shown in the bottom pane list. It will contain the same parameter values of the copied setting. The new setting can now be edited to modify any of its parameters.

Editing a setting

Select the setting to be edited from the bottom pane list. The selected setting will be highlighted. Then press the EDIT button.

When editing a setting the user is led through the same steps as when creating a new setting. If there are steps where no parameters need to be changed one can press the forward button to continue to the next step. After saving the setting the user will be redirected again to the “select parameter setting” page where all the available settings are listed.

Making a setting default

When using a particular setting very frequently the user might want to set it as default. The default setting will be automatically preselected the next time the user logs in HRM.

To make a setting default, select the setting from the bottom pane list. The selected setting will be highlighted. Press the DEFAULT button.

The name of the default setting will be highlighted in a different color when it is not selected.

Deleting a setting

To delete a setting select the setting from the bottom pane list. The selected setting will be highlighted. Press the DELETE button.

Be careful: the setting will be deleted immediately without any further questions. Once it is deleted its name will disappear from the list. A deleted setting can not be restored any more. If a setting was deleted accidentally the only way to get it back is to create a new setting and enter all values again. If this setting had been used before, the user can look for the parameter values in the summary e-mail that HRM sent when the job was finished.

Selecting the images

At this stage the images to be processed can be selected. The "available images on server" area lists all images on the user account that match the file type of the current parameter settings (see Figure 15.4). If no images are listed on this area it is quite likely that there are no images of the selected type in the user account.

SHIFT+CLICK and CONTROL+CLICK can be used to select multiple images. Press the down arrow button to add the images to the "selected images" area. Images can be removed from the selection in the same way using the up arrow button.

If more images are added to the file server after reaching the image selection stage, they will not be shown until the UPDATE VIEW button is pressed so that HRM can rescan the user account and refresh the file list.

If the image setting involved in a particular deconvolution job states that an experimental Point Spread Function (PSF) is to be used in the computation, HRM will ask for the correspoinding PSF file per channel. Thus, keep in mind that multichannel raw images require a PSF file per channel for deconvolution. For more information about the PSF, see the SVI-wiki[7].

If an image template is selected for the deconvolution of an image with more channels than what the template specifies, HRM will use theoretical PSFs for the remaining channels, up to the number of channels of the raw image to be deconvolved. Be aware that HRM does not warn about this situation.

The PSFs must be always saved in ICS format to store all the metadata necessary for Huygens Core to do the deconvolution.

The Huygens software can read plenty of file formats frequently used in fluorescence microscopy, but HRM may not be fully adapted to make use of some of these formats. For a full list of file formats supported by Huygens Core please see the most recent online list [8].

Create the job

This is the last "Start a job" page, where HRM needs to be told which output file format to use for the deconvolved images and where the deconvolution job is launched..

If the user would like to change any settings before creating the deconvolution job it is recommended to use the links "image parameters" or "restoration parameters" displayed on this page. They will take the user back to the corresponding setting selection page, where different settings can be selected or the current one be again edited.

Likewise, in order to change the images on which to run the deconvolution job one can click on the "selected images" link. This shows back the "Select images" page where the image selection can be changed.

7. http://www.svi.nl/FrontPagePointSpreadFunction
8. http://www.svi.nl/FrontPageFileFormats

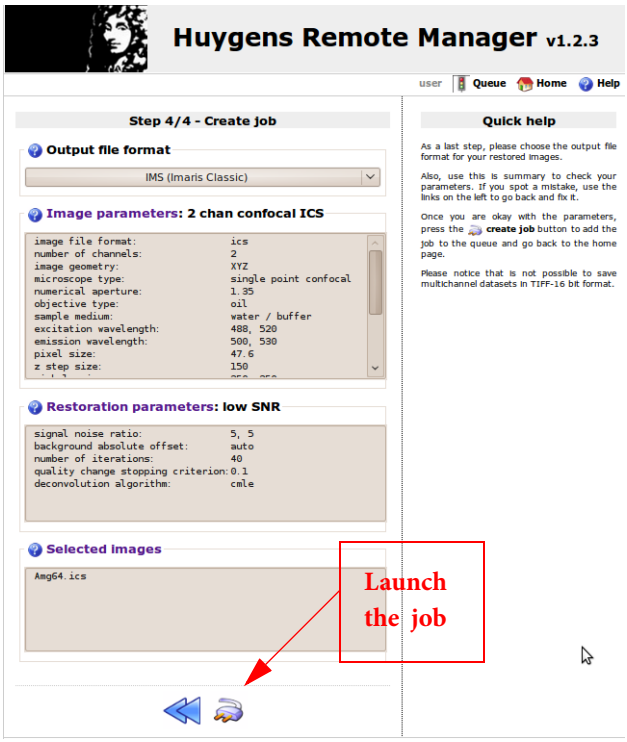

FIGURE 15.5. Reviewing the settings and creating an HRM job. The file type for the deconvolved image must be selected at this point. Cliking on the plug icon launches the deconvolution job.

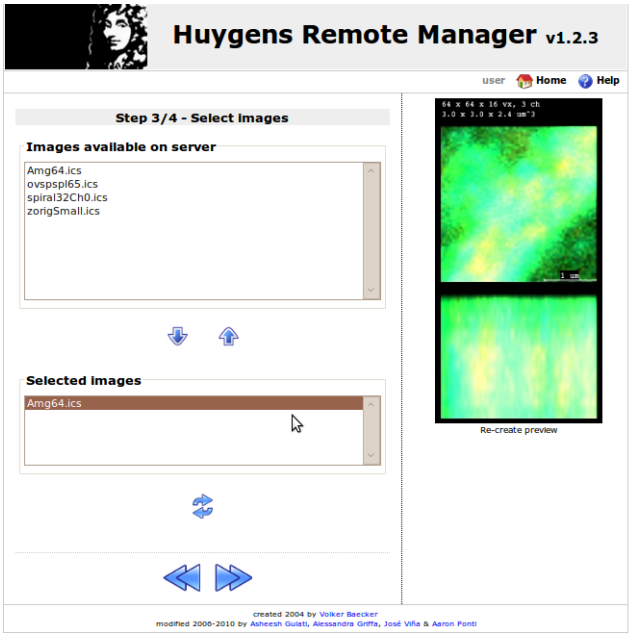

FIGURE 15.4. Image file selection in HRM. Only images matching the currently selected settings stored in the user account are shown for selection.

To create the job at this point simply click on the plug icon (see Figure 15.5). This will launch the job and will show again the HRM home page. The progress of the queued

jobs can be checked by clicking on the “queue” link at the top of the home page (see Figure 15.6).

HRM sends the user a confirmation e-mail when the deconvolution job is finished. The result images are stored in the “Results” area of the user’s HRM account.

The deconvoluton results can be checked out and downloaded using HRM. From the HRM home page, click on the File Manager button and then on the “Results” link on the top menu. There a list of deconvolved images is displayed. The images can be selected to check them out, if requested, to download them to a local machine. Uplon clicking on the blue down arrow the selected images will be packed in a single tar file and then downloaded by an Internet browser

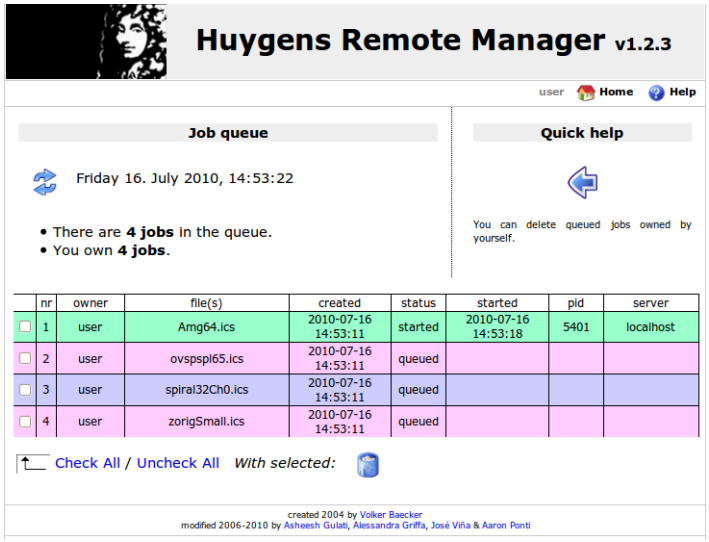

FIGURE 15.6. The HRM job queue. All the queued jobs are visible to all users, but only the jobs owned by the current user can be deleted.

CHAPTER 16

Appendix

License String Details

Detailed information about the installed license strings can be displayed via HELP→LICENSE. Select the license string of interest and click EXPLAIN LICENSE.

A Huygens license string consists of a set of substrings separated by dashes (-). These substrings describe e.g. the product, version number, options, etc. The checksum at the end of the string should match with all other substrings. A complete string looks like this:

```
HuEss-3.7-wcnp-d-tvAC-emnps-eom2012Dec31-
e7b7c623393d708e-{user@domain.com}-4fce0dbe86e8ca4344dd
```

Table 16.1 lists the building blocks from which this string is composed.

TABLE 16.1. The building blocks of the Huygens license string.

Substring	Description
Product	The product to which the license string applies. This can be `HuEss`, `HuPro`, `HuScript`, and `HuCore`.
Version	The version number of the product.
Microscope types	This substring consists of one or more characters representing the microscope types for which the deconvolution is enabled. These are '`w`' (widefield), '`c`' (confocal), '`n`' (Spinning disk), '`p`' (multi-photon), and '`4`' (4-Pi experimental microscopes).
Server flag	Determines the number of cores that are enabled for multi-threading. A hyperthreaded core is counted as a single core. It can be '`d`' (desktop; 2 cores), '`s`' (small server; 4 cores), '`m`' (medium server; 8 cores), '`l` (larger server; 16 cores), and '`x`' (extreme server; 512 cores).
Option flags	This is a set of characters that list the enabled optional modules. An overview of these modules is given in Table 16.2.
Locking policy	A set of characters that indicate to which properties the license is locked. These can be '`d`' (expiry date), '`e`' (e-mail address), '`m`' (system ID), '`n`' (number of cores), '`p`' (processor type), and '`s`' (processor details).

TABLE 16.1. The building blocks of the Huygens license string.

Substring	Description
Expiry date	The date on which the license or maintenance ends. When this substring starts with 'eom', then only the maintenance expires; the license remains valid.
System ID	A 16 character hexadecimal string containing hardware identification numbers.
E-mail address	The customer e-mail address.
Checksum	A 20 character hexadecimal checksum on the previous substrings.

TABLE 16.2. Identifiers for the optional modules.

Character	Module
a	Object Analyzer
A	Advanced Object Analyzer
b	Small file reader bundle
B	Complete file reader bundle
C	Colocalization Analyzer
f	PSF Distiller
M	Movie Maker
N	Enable new, pre-released features
s	Object Stabilizer
t	Time series
v	Surface Renderer visualization

The Point Spread Function

One of the basic concepts in image deconvolution is the *point spread function* (PSF). The PSF of the microscope is the image which results from imaging a point object in the microscope. Because of wave diffraction[1] a point object is imaged (spread out) into a fuzzy spot: the point spread function. In fluorescence imaging the PSF completely determines the image formation. In other words: *all microscopic imaging properties are packed into this 3D function*. In Huygens Essential, a PSF can be obtained in two different ways:

1. *Generating a theoretical PSF*: When a measured PSF is not available, Huygens Essential automatically uses a theoretical PSF. The PSF is computed from the microscopic parameters attached to the data. Because a theoretical PSF can be generated without any user intervention Huygens Essential does the calculation in the background without any notice.
   Images affected by spherical aberration (See "Refractive Index Mismatch" on page 117) are better restored using a theoretical depth-dependent PSF.
2. *Measuring a PSF*: By using the PSF Distiller a measured PSF can be derived from images of small fluorescent beads (See "Beads Suited for PSF Distillation" on page 108). Measured PSF's improve deconvolution results and may also serve as a quality test for the microscope

1. http://www.svi.nl/FrontPageImageFormation

Quality Factor

Deconvolution as it is done in Huygens Essential is based on the idea of finding the best estimate of the object that is imaged by the microscope. To assess the quality of an estimate, Huygens Essential simulates the microscopic imaging of each estimate (the estimated is convolved with the PSF) and compares the simulation with the measured image. From the difference a quality factor is computed. The difference is also used to compute a correction factor to modify the estimate in such a way that the corrected estimate will yield a better quality factor. The quality factor as reported by the software is a measure relative to the first estimate and therefore a number greater than or equal to 1. If the increase in quality drops below the *quality threshold* the iterations are stopped.

File Series

There are many ways in which Tiff files or other file series are named. These files can have multiple counters (referring to *slices*, *time frames*, or *channels*), and these counters can have arbitrary prefixes and ordering.

Numbered Tiff Series

If a series is simply numbered like: `slice001.tif`, `slice002.tif`, `...`, `slice0nn.tif`, then Huygens Essential will read the series into a single 3D image. Because Tiff files usually carry no additional microscopic information, check the parameters carefully.

Leica Numbering

Huygens Professional natively supports both reading and *writing* Tiff series with Leica style numbering, if there is more than one channel, slice, or time frame. A single channel 2D time series would be numbered according to the scheme:

```
im_tNN.tif
```

Here, `NN` is replaced by the time index for each frame. A more complex, multi-channel 3D time series has this pattern:

```
im_tNN_zNNN_cNN.tif
```

In this series, the second channel of the fourth slice of the third time frame has the filename:

```
im_t02_z003_c01.tif
```

The File Series Tool

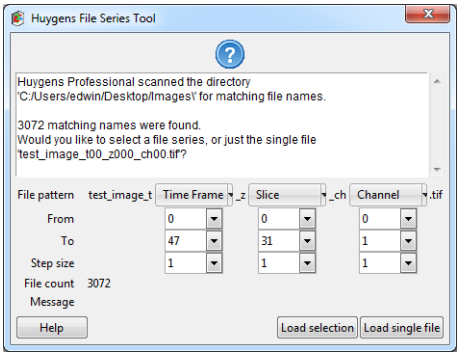

FIGURE 16.1. The Huygens File Series Tool automatically scans a directory for a file series.

Although Huygens Essential uses Leica style numbering for writing files, the software attempts to detect *any type* of file series for reading. Whenever a file is opened that appears to be part of a file series, Huygens Essential shows the *File Series Tool* dialog (Figure 16.1). This tool enables the user to select a subset of a file series, and select a dimension for each the indices in the file name, so that each image is assigned to the correct *z*-plane, time frame, and channel.

The file pattern is shown in the first row in the dialog. The counters in the file name are replaced by menu buttons for selecting the appropriate dimension for each counter. The options are:

- **Slice**: The range of this counter becomes the *z*-dimension.
- **Time Frame**: The range of this counter becomes the time dimension.
- **Channel**: The range of this counter becomes the channel dimension.
- **Ignore**: the variable is ignored. This is useful to omit e.g. the value of time stamps.
- **The value of the counter** in the selected file: the value of this counter has to match the value in the selected file.

Note that the selection has to be unique, i.e. it is impossible to have ignored variables without having a *Slice*, *Time Frame*, or *Channel* counter.

In the second, third, and fourth row, the range for each of the counters can be defined. A range from 0 to 9 with step size 2 will load the files 0, 2, 4, 6, and 8. Note that the time (in seconds) and *z*-sampling intervals (in nm) are not adapted to the step sizes.

Press the LOAD SELECTION button to load all files in the series into a single image. Before the dialog is closed, the tool will check if all files in the selection are really present in the directory.

Adjusting the Global Color Scheme

Huygens Essential uses a global scheme for coloring the different channels in multi-channel images. These colors can be adjusted through the *Preferences* window via TOOLS→PREFERENCES...→EDIT GLOBAL COLORS (See Figure 16.2).

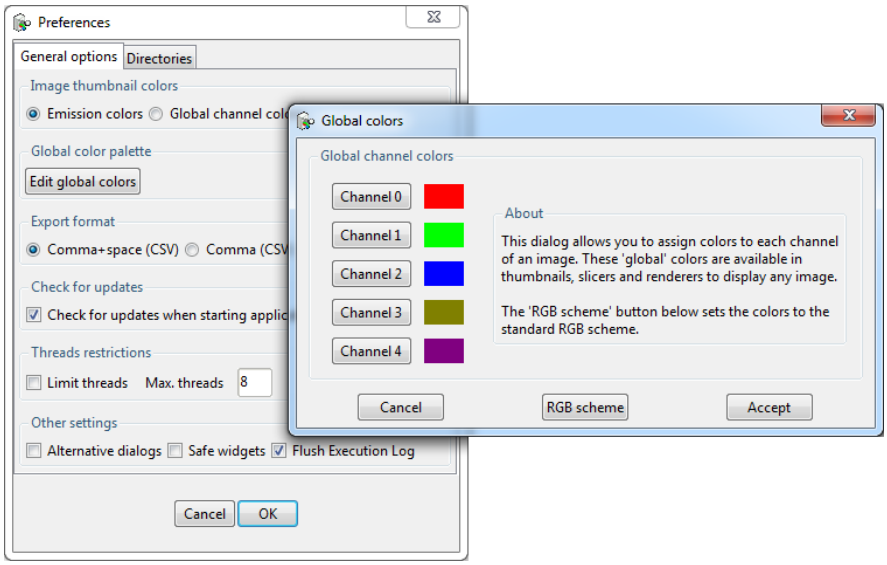

FIGURE 16.2. The global color scheme can be modified through the *Preferences* window.

Hue Selector

The *hue selector* is a component that allows adjustment of the color range in which objects are displayed (See "The Surface Renderer" on page 37 and "The Colocalization Analyzer" on page 81). Objects belonging to different channels can be represented in different hue ranges to make them clearly distinct. The gradual differences inside the selected range make independent objects distinguishable. Also a range can be collapsed to have all objects in a channel displayed with exactly the same color. In Huygens the *hue selector* does appear in two flavors.

Hue Range

This selector allows the adjustment of a *hue* range. The objects on which this selector acts will get a color that lies within this range. The assignment of colors is based on the position of an object or on another parameter.

Hue Range and Saturation

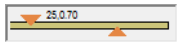

This selector allows the adjustment of a single *hue* value and a *saturation*. The upper triangle defines the color, while the lower triangle sets the saturation for this color; left is white, right is fully saturated.

Image Statistics

Right-click on a thumbnail image and select SHOW PARAMETERS from the pop-up menu. This window shows, besides the parameter settings, statistical information of the particular image. Amongst them are the mean, sum, standard deviation, norm, and position of the center of mass.

Setting the Coverslip Position

When there is a mismatch between the refractive index for which the microscope's objective is designed and the actual refractive index of the embedding medium, the shape of the point spread function (PSF) will be distorted due to spherical aberration (See "Refractive Index Mismatch" on page 117). As deeper layers in the specimen are imaged, moving away from the coverslip, this distortion will progressively worsen. To compute the spherical aberration it is necessary to know the distance from the coverslip. Because in many cases the coverslip position does not coincide with the first plane in the data, this position can be set in the microscopic parameter editor. To our knowledge none of the existing microscopic image files record the coverslip position in the meta data.

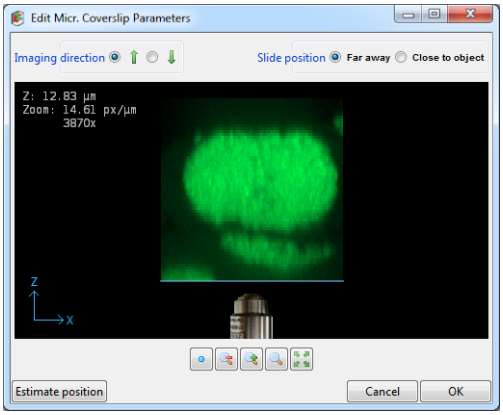

FIGURE 16.3. The coverslip position editor showing an *xy* MIP of the data along they *y*-direction. The coverslip position can be adjusted by dragging the blue line. The imaging direction, here *upwards*, is indicated by the position of the objective relative to the data as shown. The *z*-position shown top-left in the image indicates the distance in µm of the coverslip to the first data plane.

Next to direct numerical input, the coverslip position and imaging direction can be set using a visual editor (Figure 16.3), reachable from the parameter editor or from the Huygens Essential wizard by clicking the wrench button ( ).

Inverted Microscope

The editor shows the coverslip position and imaging direction relative to the data as read from the microscopic file. In an inverted microscope, with the objective physically below the specimen it is likely that the first *xy*-plane in the data, corresponding with the lowest location in the *xz* maximum intensity projection (MIP) on the screen, corresponds with the *xy*-plane scanned closest to the objective. This situation is shown in REF TO FIG. However, since scan directions and data planes might have been reordered, this match is not guaranteed. Fortunately, it is often easy to spot the flat side of the object where it adheres to the glass, so the orientation can be verified.

Upright Microscope

In an upright microscope, and a *z*-scan starting away from the coverslip, the first plane is also likely to be physically the lowest plane. In that case, the imaging direction should be set to downwards and the coverslip position in the top part of the *xz* MIP projection. However, if the scan started close to the coverslip while storing these first planes first in the data set, the MIP projection will show the data upside down. Consequently, the coverslip position will be in the lower part of the MIP, and the imaging direction is upward.

Slide Position

When the specimen is mounted on the coverslip, the distance from the object to the slide is probably in the range from 50 to100 µm, outside of the image. In this case, or in the case there is no slide, select *Far away* in the top-right selector.

When the specimen is close to or mounted on the slide, select *Close to object* (). Drag the coverslip to its proper location. When this location is at some distance from the data it might be necessary to zoom out. The image can be dragged by holding down the right mouse button. In terms of imaging quality, when there is a refractive index mismatch between embedding medium and immersion medium, this is not an ideal situation since the light from and to the objective must travel hundreds of wavelengths through the embedding medium, possibly resulting in strong spherical aberration induced bloating of the PSF.

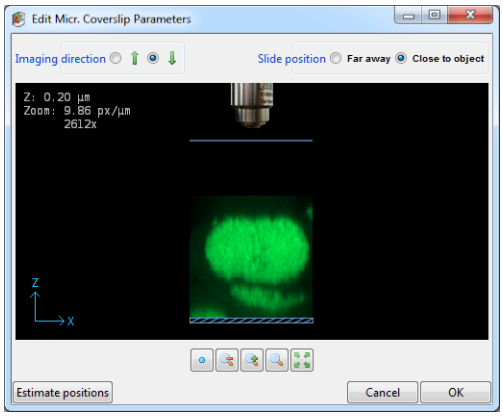

FIGURE 16.4. The Coverslip editor with the slide position set to *Close to object*.

Excitation Beam Overfill Factor

In confocal microscopes, the entry pupil of the microscope objective is illuminated by a laser beam. Usually, laser beams have a Gaussian intensity profile[2]. As a result, the illumination intensity is not constant over the pupil but will decrease towards the edges. Lower edge intensities will lower the effective NA and therefore negatively affect resolution. In most confocal microscopes this is remedied by using a beam width which is significantly larger than the entry pupil, at the cost of loss of excitation power. The ratio between the beam width and the pupil diameter is the excitation beam overfill factor (See Figure 16.5) and is typically in the range from 2 to 4.

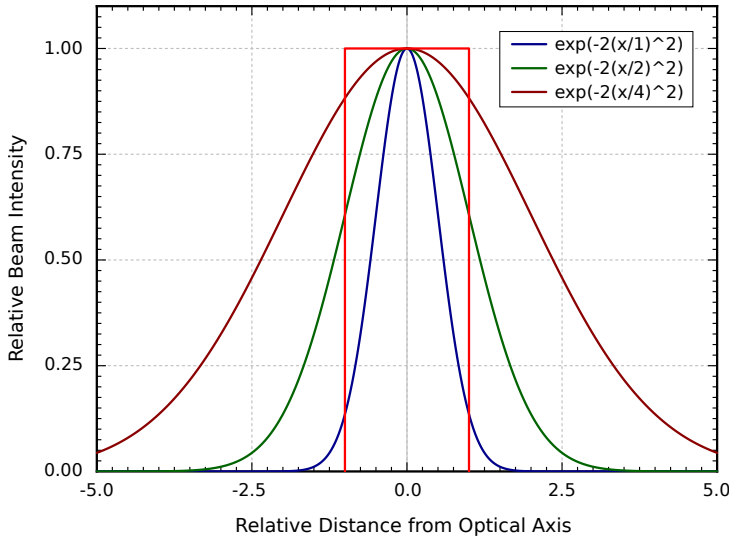

FIGURE 16.5. Lens entry pupil (*red*), beam profile with overfill factor 1 (*blue*), 2 (*green*), and 4(*dark red*). At overfill factor 1, the beam intensity is 14 % of the maximum, at overfill factor 2 the edge intensity is 61 % of the maximum.

2. http://en.wikipedia.org/wiki/Gaussian_beam

The overfill factor can be set as a microscopic parameter in Huygens Essential, and is taken into account when computing the point spread function.

Brightfield Images

Brightfield imaging is not a *linear imaging* process. In a linear imaging process the image formation can be described as the linear convolution of the object distribution and the point spread function, hence the name deconvolution for the reverse process. So in principle one cannot apply deconvolution based on linear imaging to non linear imaging modes like brightfield and reflection. One could state that the image formation in these cases *is* linear because it is governed by linear superposition of amplitudes. However, microscopes do not measure light amplitudes but rather intensities, i.e. the absolute squared values of the amplitudes. Taking the absolute square destroys all phase information one would need to effectively apply deconvolution. Fortunately, in the brightfield case the detected light is to a significant degree incoherent. Because in that case there are few phase relations the image formation is largely governed by the addition of intensities, especially if one is dealing with a high contrast image.

In practice one goes about deconvolving brightfield images by inverting them (using TOOLS→INVERT IMAGE image) and processing them further as incoherent fluorescence widefield images. The Tikhonov Miller algorithm was proven to work excellently for brightfield data. This algorithm is available in the Huygens Professional only. With the MLE algorithm one should watch out sharply for interference like patterns (periodic rings and fringes around objects) in the measured image. As a rule these become pronounced in low contrast images. After the deconvolution run a reverse to the original contrast setting is possible.

Support and Contact Information

Addresses and Phone Numbers

Mailing Address	Scientific Volume Imaging B.V. Laapersveld 63 1213 VB Hilversum The Netherlands
Phone	+31 35 6421626
Fax	+31 35 6837971
E-mail	`info@svi.nl`
URL	`http://www.svi.nl`

We are directly reachable by phone during office hours (CET) or by e-mail 24/7.

Distributors

A up-to-date list of distributors can be found on our web site[3].

SVI Support Wiki

The SVI-wiki[4] is a rapidly expanding public knowledge resource on 3D microscopy and deconvolution. Based on the WikiWikiWeb principle, it is open to contributions from

3. `http://www.svi.nl/distributors/`
4. `http://www.svi.nl/FrontPage`

every visitor. In addition it serves as a support medium for SVI customers and relations to discuss different aspects of the Huygens software.

This is a list of useful starting points in the SVI Support Wiki to learn more about the Huygens software and microscopical imaging in general:

- Information on the parameters describing the imaging conditions (sampling, numerical aperture, pinholes, etc.):
  `http://support.svi.nl/wiki/MicroscopicParameters`
- Information on the restoration parameters (signal to noise ratio, background, quality criteria, etc.) used by the deconvolution algorithms:
  `http://support.svi.nl/wiki/RestorationParameters`
- A step by step example on how to tune these parameters to achieve the desired restoration results:
  `http://support.svi.nl/wiki/DeconvolutionProcedure`
- Important issues regarding image acquisition and restoration (sampling, clipping, etc.):
  `http://support.svi.nl/wiki/ImportantFactors`
- Typical acquisition pitfalls (spherical aberration, undersampling, bleaching, etc.):
  `http://support.svi.nl/wiki/AcquisitionPitfalls`
- Information on recording beads to measure a PSF:
  `http://support.svi.nl/wiki/RecordingBeads`
- Tutorials and detailed information on using the different aspects of the Huygens software (restoration, visualization, analysis, programming, etc.):
  `http://support.svi.nl/wiki/Tutorials`
- Uploading images to SVI:
  `http://support.svi.nl/wiki/SendImagesToSvi`

A
Address 4, 108
Analyze object 53
Anchor 72
Animation
  frame count 31, 36, 40
  frame rate 31, 36, 40
  MIP 31
  Movie Maker 43
  SFP 36
  Surface Renderer 40
AVI files 46

B
Background 14, 84
Backprojected
  pinhole radius 10
  pinhole spacing 10
Black level 14
Bleaching correction 15
Bounce 48
Brick layout 15
Brightfield images 108
Brightness 21

C
Center scene 31, 40
Channels 16
Clipping 14
CMLE 15
Colocalization 63, 81
  coefficients 83
  map 83
Color mode 31, 36
Colors 21, 22, 104
Compress 21
Contact SVI 108
Contrast 21
Convert 10
Correlation 63
Coverslip position 10, 106
Cropper 12

D
Deconvolution
  wizard 9
Discard object 53
Distributors 108
Drift 17

E
E-mail 4, 108
Emission
  transparency 35
  wavelength 11
Excitation
  fill factor 11, 107
  photon count 11
  transparency 35
  wavelength 11
Experiment preset 56, 73

F
Fax number 4, 108
File
  Formats 9
  Open 9
  Save 17
  Series 103
Fill factor 11
Filter Objects 75
Freeware mode 4

G
Gamma 21
Garbage volume 39, 52
Global color scheme 104

H
hgst file 12
hgsv file 31, 36, 40
Histogram 13, 58

I
Imaging direction 10
Installation 3
  Linux 3
  Mac OS X 3
  path 6
  Uninstall 7
  Windows 3
Intensity profile 24
Iso-surface 37
Iteration mode 15
Iterations 15

K
Keyframes 44

L
Leica numbering 103
Lens refractive index 10
License string 4, 101
Linking 23
Linux 3

M
Mac OS X 3
Markers 24
Maximum Intensity Projection 29
Measure
  Aspect Ratio 63
  Colocalization 63, 81
  Correlation 63
  Intensity profile 24
  Length 24, 61
  Principal Axis 61
  Sphericity 62
Medium refractive index 10
Memory 7
Microscope type 10
Microscopic parameters 10
MIP 29

Mouse mode 53, 68
Movie Maker 43
Multi-channel images 16

N
Number of iterations 15
Numerical aperture 10

O
Object Analyzer 49, 87
Object size 34
Objective quality 10
Orthogonal Slicer 27
Overfill factor 107

P
Parameter templates 12
Parameters 10
Phone number 4, 108
Pinhole
  radius 10
  spacing 10
Pipe
  Mode 70
  Render 38, 54, 74
Point spread function 102
Presets 46, 56, 73
Principal axis 61
Processor 7
PSF 102

Q
Quality factor 103
Quality threshold 15

R
Refractive index
  Lens 10
  Medium 10
Region of interest 66
Render
  MIP 29
  pipes 38, 54, 74
  quality 31, 36
  SFP 33
  size 31, 36, 40
  Surface 37
ROI 66

S
Saturation 14
Segmentation 51
Select area 53
Select object 53
SFP 33
Signal to noise ratio 15
Slide position 106
SNR 15
Spherical aberration 106
Sphericity 62
Statistics 55, 105
Storyboard 45
Support 108
Surface Renderer 37
System ID 4, 102
System requirements 7

T
Table style 36
Templates
  Microscopic 12
  Visualization 31, 36, 40
Threshold 29, 34, 38, 51
Time series 10
Timeline 46
Transitions 45
Transparency 39
Transparency depth 40
Twin Slicer 19

U
Uninstall 7
Updates 5, 6

W
Wavelength
  Emission 11
  Excitation 11
Windows 3
Wizard 9